



The NANOGrav 15 yr Data Set: Constraints on Supermassive Black Hole Binaries from the Gravitational-wave Background

Gabriella Agazie¹ , Akash Anumalapurdi¹ , Anne M. Archibald² , Paul T. Baker³ , Bence Bécsey⁴ , Laura Blecha⁵ , Alexander Bonilla^{6,7} , Adam Brazier^{8,9} , Paul R. Brook¹⁰ , Sarah Burke-Spolaor^{11,12} , Rand Burnette⁴ , Robin Case⁴ , J. Andrew Casey-Clyde¹³ , Maria Charisi¹⁴ , Shami Chatterjee⁸ , Katerina Chatziioannou¹⁵ , Belinda D. Cheeseboro^{11,12} , Siyuan Chen¹⁶ , Tyler Cohen¹⁷ , James M. Cordes⁸ , Neil J. Cornish¹⁸ , Fronefield Crawford¹⁹ , H. Thankful Cromartie^{8,71} , Kathryn Crowter²⁰ , Curt J. Cutler^{15,21} , Daniel J. D’Orazio²² , Megan E. DeCesar²³ , Dallas DeGan⁴ , Paul B. Demorest²⁴ , Heling Deng⁴ , Timothy Dolch^{25,26} , Brendan Drachler^{27,28} , Elizabeth C. Ferrara^{29,30,31} , William Fiore^{11,12} , Emmanuel Fonseca^{11,12} , Gabriel E. Freedman¹ , Emiko Gardiner³² , Nate Garver-Daniels^{11,12} , Peter A. Gentile^{11,12} , Kyle A. Gersbach¹⁴ , Joseph Glaser^{11,12,74} , Deborah C. Good^{13,33} , Kayhan Gültekin³⁴ , Jeffrey S. Hazboun⁴ , Sophie Hourihane¹⁵ , Kristina Islo¹ , Ross J. Jennings^{11,12,71} , Aaron Johnson^{1,35} , Megan L. Jones¹ , Andrew R. Kaiser^{11,12} , David L. Kaplan¹ , Luke Zoltan Kelley³² , Matthew Kerr³⁶ , Joey S. Key³⁷ , Nima Laal⁴ , Michael T. Lam^{27,28} , William G. Lamb¹⁴ , T. Joseph W. Lazio²¹ , Natalia Lewandowska³⁸ , Tyson B. Littenberg³⁹ , Tingting Liu^{11,12} , Jing Luo^{40,73} , Ryan S. Lynch⁴¹ , Chung-Pei Ma^{32,42} , Dustin R. Madison⁴³ , Alexander McEwen¹ , James W. McKee^{44,45} , Maura A. McLaughlin^{11,12} , Natasha McMann¹⁴ , Bradley W. Meyers^{20,46} , Patrick M. Meyers¹⁵ , Chiara M. F. Mingarelli^{13,33,47} , Andrea Mitridate⁴⁸ , Priyamvada Natarajan^{49,50} , Cherry Ng⁵¹ , David J. Nice⁵² , Stella Koch Ocker⁸ , Ken D. Olum⁵³ , Timothy T. Pennucci⁵⁴ , Benetge B. P. Perera⁵⁵ , Polina Petrov¹⁴ , Nihan S. Pol¹⁴ , Henri A. Radovan⁵⁶ , Scott M. Ransom⁵⁷ , Paul S. Ray³⁶ , Joseph D. Romano⁵⁸ , Jessie C. Runnoe¹⁴ , Shashwat C. Sardesai¹ , Ann Schmiedekamp⁵⁹ , Carl Schmiedekamp⁵⁹ , Kai Schmitz⁶⁰ , Levi Schult¹⁴ , Brent J. Shapiro-Albert^{11,12,75} , Xavier Siemens^{1,4} , Joseph Simon^{61,72} , Magdalena S. Siwek⁶² , Ingrid H. Stairs²⁰ , Daniel R. Stinebring⁶³ , Kevin Stovall²⁴ , Jerry P. Sun⁴ , Abhimanyu Susobhanan¹ , Joseph K. Swiggum^{52,72} , Jacob Taylor⁴ , Stephen R. Taylor¹⁴ , Jacob E. Turner^{11,12} , Caner Unal^{64,65} , Michele Vallisneri^{21,35} , Sarah J. Vigeland¹ , Jeremy M. Wichter⁶⁶ , Haley M. Wahl^{11,12} , Qiaohong Wang¹⁴ , Caitlin A. Witt^{12,67,68} , David Wright⁶⁹ , and Olivia Young^{27,28}

The NANOGrav Collaboration⁷⁰

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

² Newcastle University, NE1 7RU, UK

³ Department of Physics and Astronomy, Widener University, One University Place, Chester, PA 19013, USA

⁴ Department of Physics, Oregon State University, Corvallis, OR 97331, USA

⁵ Physics Department, University of Florida, Gainesville, FL 32611, USA

⁶ Observatório Nacional, Rua General José Cristino 77, São Cristóvão, 20921-400 Rio de Janeiro, RJ, Brazil

⁷ Facultad de Ciencias, Departamento de Matemáticas, Universidad El Bosque, Ak. 9 # 131 A-2, Bogotá, Colombia

⁸ Cornell Center for Astrophysics and Planetary Science and Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

⁹ Cornell Center for Advanced Computing, Cornell University, Ithaca, NY 14853, USA

¹⁰ Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

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

¹² Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

¹³ Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA

¹⁴ Department of Physics and Astronomy, Vanderbilt University, 2301 Vanderbilt Place, Nashville, TN 37235, USA

¹⁵ Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

¹⁶ Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing, 100871 People’s Republic of China

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

²² Niels Bohr International Academy, Niels Bohr Institute, Blegdamsvej 17, DK-2100 Copenhagen, Denmark

²³ George Mason University, Resident at the Naval Research Laboratory, Washington, DC 20375, USA

²⁴ National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM 87801, USA

²⁵ Department of Physics, Hillsdale College, 33 East College Street, Hillsdale, MI 49242, USA

²⁶ Eureka Scientific, 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA

²⁷ School of Physics and Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA

²⁸ Laboratory for Multiwavelength Astrophysics, Rochester Institute of Technology, Rochester, NY 14623, USA

²⁹ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

³⁰ Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771

³¹ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

³² Department of Astronomy, University of California, Berkeley, 501 Campbell Hall #3411, Berkeley, CA 94720, USA

³³ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

³⁴ Department of Astronomy and Astrophysics, University of Michigan, Ann Arbor, MI 48109, USA

³⁵ Theoretical AstroPhysics Including Relativity (TAPIR), MC 350-17, California Institute of Technology, Pasadena, CA 91125, USA

³⁶ Space Science Division, Naval Research Laboratory, Washington, DC 20375-5352, USA

³⁷ University of Washington Bothell, 18115 Campus Way NE, Bothell, WA 98011, USA

³⁸ Department of Physics, State University of New York at Oswego, Oswego, NY 13126, USA

³⁹ NASA Marshall Space Flight Center, Huntsville, AL 35812, USA

- ⁴⁰ Department of Astronomy & Astrophysics, University of Toronto, 50 Saint George Street, Toronto, ON M5S 3H4, Canada
⁴¹ Green Bank Observatory, P.O. Box 2, Green Bank, WV 24944, USA
⁴² Department of Physics, University of California, Berkeley, CA 94720, USA
⁴³ Department of Physics, University of the Pacific, 3601 Pacific Avenue, Stockton, CA 95211, USA
⁴⁴ E.A. Milne Centre for Astrophysics, University of Hull, Cottingham Road, Kingston-upon-Hull, HU6 7RX, UK
⁴⁵ Centre of Excellence for Data Science, Artificial Intelligence and Modelling (DAIM), University of Hull, Cottingham Road, Kingston-upon-Hull, HU6 7RX, UK
⁴⁶ International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia
⁴⁷ Department of Physics, Yale University, New Haven, CT 06520, USA
⁴⁸ Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, D-22607 Hamburg, Germany
⁴⁹ Department of Astronomy, Yale University, 52 Hillhouse Avenue, New Haven, CT 06511, USA
⁵⁰ Black Hole Initiative, Harvard University, 20 Garden Street, Cambridge, MA 02138, USA
⁵¹ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 Saint George Street, Toronto, ON M5S 3H4, Canada
⁵² Department of Physics, Lafayette College, Easton, PA 18042, 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, University of Puerto Rico, Mayagüez, PR 00681, USA
⁵⁷ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA
⁵⁸ Department of Physics, Texas Tech University, Box 41051, Lubbock, TX 79409, USA
⁵⁹ Department of Physics, Penn State Abington, Abington, PA 19001, USA
⁶⁰ Institute for Theoretical Physics, University of Münster, D-48149 Münster, Germany
⁶¹ Department of Astrophysical and Planetary Sciences, University of Colorado, Boulder, CO 80309, USA
⁶² Center for Astrophysics, Harvard University, 60 Garden Street, Cambridge, MA 02138, USA
⁶³ Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA
⁶⁴ Department of Physics, Ben-Gurion University of the Negev, Be'er Sheva 84105, Israel
⁶⁵ Feza Gursey Institute, Bogazici University, Kandilli, 34684, Istanbul, Turkey
⁶⁶ Department of Physics, Skidmore College, 815 North Broadway, Saratoga Springs, NY 12866, USA
⁶⁷ Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208, USA
⁶⁸ Adler Planetarium, 1300 South DuSable Lake Shore Drive, Chicago, IL 60605, USA
⁶⁹ Department of Physics, University of Central Florida, Orlando, FL 32816-2385, USA
- Received 2023 June 13; revised 2023 June 22; accepted 2023 June 26; published 2023 August 1

Abstract

The NANOGrav 15 yr data set shows evidence for the presence of a low-frequency gravitational-wave background (GWB). While many physical processes can source such low-frequency gravitational waves, here we analyze the signal as coming from a population of supermassive black hole (SMBH) binaries distributed throughout the Universe. We show that astrophysically motivated models of SMBH binary populations are able to reproduce both the amplitude and shape of the observed low-frequency gravitational-wave spectrum. While multiple model variations are able to reproduce the GWB spectrum at our current measurement precision, our results highlight the importance of accurately modeling binary evolution for producing realistic GWB spectra. Additionally, while reasonable parameters are able to reproduce the 15 yr observations, the implied GWB amplitude necessitates either a large number of parameters to be at the edges of expected values or a small number of parameters to be notably different from standard expectations. While we are not yet able to definitively establish the origin of the inferred GWB signal, the consistency of the signal with astrophysical expectations offers a tantalizing prospect for confirming that SMBH binaries are able to form, reach subparsec separations, and eventually coalesce. As the significance grows over time, higher-order features of the GWB spectrum will definitively determine the nature of the GWB and allow for novel constraints on SMBH populations.

Unified Astronomy Thesaurus concepts: [Gravitational waves \(678\)](#); [Supermassive black holes \(1663\)](#); [Galaxy evolution \(594\)](#)

1. Introduction

Strong observational evidence suggests that most, if not all, massive galaxies contain supermassive black holes (SMBHs) at their centers (Richstone et al. 1998). Additionally, hierarchical structure formation causes frequent galaxy mergers (Ostriker &

Hausman 1977; White 1980; Lacey & Cole 1993), naturally leading to the formation of SMBH binaries, which may also merge (Begelman et al. 1980; Milosavljević & Merritt 2001). At the last stages of their evolution, these binaries produce strong nanohertz gravitational-wave (GW) emission that can be targeted by pulsar timing arrays (PTAs), which systematically monitor a large number of millisecond pulsars. By detecting coherent deviations in the times of arrival (TOAs) of pulsar signals, PTAs can observe a stochastic gravitational-wave background (GWB) from the superposition of many unresolved binaries, as well as individually resolved sources on top of the background (Burke-Spolaor et al. 2019; Taylor 2021).

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) 12.5 yr data set showed evidence of a common-spectrum red-noise process consistent with a GWB (Arzoumanian et al. 2020). This result was confirmed by

⁷⁰ comments@nanograv.org

⁷¹ NASA Hubble Fellowship: Einstein Postdoctoral Fellow.

⁷² NANOGrav Physics Frontiers Center Postdoctoral Fellow.

⁷³ Deceased.

⁷⁴ NSF Astronomy and Astrophysics Postdoctoral Fellow.

⁷⁵ Giant Army, 915A 17th Avenue, Seattle, WA 98122.



the Parkes Pulsar Timing Array (PPTA; Goncharov et al. 2021), the European Pulsar Timing Array (Chen et al. 2021), and the International Pulsar Timing Array (IPTA; Antoniadis et al. 2022). The NANOGrav 15 yr data set shows that the common uncorrelated red-noise (CURN) signal discovered in Arzoumanian et al. (2020) persists with greater significance and is now detected in a larger number of pulsars (Agazie et al. 2023a, hereafter *NG15gwb*). Additionally, for the first time, there is evidence of interpulsar correlations following the characteristic Hellings–Downs (HD) pattern (Hellings & Downs 1983) expected for an isotropic GWB. Careful analyses of the detection significance give false-alarm probabilities of $\approx 10^{-4}$ – 10^{-3} ($\approx 3\sigma$).

In this paper, we investigate whether the NANOGrav 15 yr results can be explained as a stochastic GWB produced by a cosmic population of SMBH binaries. While SMBH binaries have long been expected to produce such a background, a wide variety of alternative models exist, many of which invoke new physics that departs from the standard model and Λ cold dark matter (Λ CDM) cosmology. We refer the reader to Afzal et al. (2023, hereafter *NG15newphys*) for an analysis of the NANOGrav 15 yr results in the context of new-physics models such as cosmic inflation, scalar-induced GWs, domain walls, cosmic strings, and first-order phase transitions.

1.1. The Galaxy–SMBH Connection

Our understanding of galaxy formation and evolution has rapidly progressed in the last few decades. This includes the definitive and now direct observation of SMBHs in galaxy centers (Ghez et al. 1998; GRAVITY Collaboration et al. 2018; Event Horizon Telescope Collaboration et al. 2019, 2022). The mass of the central SMBH strongly correlates with the global properties of the host galaxy (e.g., the stellar velocity dispersion of the galactic bulge and the bulge mass and luminosity), with tight correlations spanning several orders of magnitude in SMBH mass (Dressler 1989; Kormendy 1993; Magorrian et al. 1998; Gebhardt et al. 2000; Tremaine et al. 2002; Häring & Rix 2004; Gültekin et al. 2009a; Kormendy & Ho 2013; McConnell & Ma 2013; Saglia et al. 2016). These trends strongly imply coordinated evolution between SMBHs and their host galaxies, which may be driven by a variety of mechanisms such as galaxy mergers, secular dynamics, stellar feedback, and feedback from active galactic nuclei (AGN; Di Matteo et al. 2005; Hopkins et al. 2008; Somerville et al. 2008). The SMBHs are believed to play particularly significant roles in shaping the structure of massive galaxies (Croton et al. 2006; Fabian 2012; Vogelsberger et al. 2014; Schaye et al. 2015; Weinberger et al. 2017), but many fundamental aspects, such as the formation channels of SMBH seeds in the early Universe or how AGN feedback shapes the host galaxies, are still poorly constrained via observations. The relevant physical processes are also very difficult to model theoretically, as they span size scales from galaxies (~ 10 kpc) to SMBH event horizons ($\sim 10^{-5}$ pc). Similar challenges limit our ability to directly model the process of SMBH binary formation and evolution.

1.2. SMBH Binary Evolution

The formation of SMBH binaries begins with the merger of two galaxies, each hosting a central SMBH. At different stages of the evolution of the SMBH pair, different physical processes dominate energy and angular momentum extraction, which

drives the binary to closer separations (Begelman et al. 1980; see De Rosa et al. 2019 for a recent review). Initially, the SMBHs are a gravitationally unbound pair (a dual SMBH) falling toward the center of the merging host (Barnes & Hernquist 1992) via dissipative “hardening” processes, such as dynamical friction (Chandrasekhar 1943; Antonini & Merritt 2012). Once the mass enclosed within the orbit is comparable to the mass of the binary (typically at \sim parsec-scale separations), the two black holes become a gravitationally bound pair (an SMBH binary; Merritt & Milosavljević 2005). At these separations, the timescale for the GW-driven inspiral is generally still longer than the Hubble time, and their GW frequencies are orders of magnitude below those that PTAs can probe.

The astrophysical environment of the binary is therefore crucial for bringing these systems to the PTA bands and, ultimately, their final coalescence. Scattering of individual stars that pass close to the SMBHs can extract energy and angular momentum from the system, hardening the binary orbit (Yu 2002). In some cases, the supply of stars on close orbits may be insufficient, and the binary will fail to merge within a Hubble time (Begelman et al. 1980). However, this so-called “final-parsec problem” has a number of potential theoretical solutions (e.g., Berczik et al. 2006; Holley-Bockelmann & Sigurdsson 2006; Khan et al. 2011; Holley-Bockelmann & Khan 2015). Similarly, in gas-rich systems, circumbinary gas disks can also catalyze the binary evolution (Escala et al. 2005; Dotti et al. 2007; Haiman et al. 2009), but the efficiency of this process or whether the gas pushes the binary inward or outward is still unclear (Moody et al. 2019; Muñoz et al. 2019; Duffell et al. 2020; Siwek et al. 2023).

If a binary stalls for longer than the time between successive galaxy mergers, a second galaxy could bring a third SMBH into the system. Triple SMBH interactions can greatly reduce the timescale for an SMBH binary merger and may also cause the ejection of the lightest SMBH from the system (Saslaw et al. 1974; Volonteri et al. 2003; Hoffman & Loeb 2007; Bonetti et al. 2016, 2018a). Once an SMBH binary reaches a sufficiently small separation, the GWs will dominate its evolution, carrying away energy and angular momentum and leading the SMBHs to coalescence (Peters & Mathews 1963).

1.3. Electromagnetic Signatures of SMBH Binaries and Multimessenger Prospects

Many studies have used electromagnetic observations of AGN to find candidate SMBH pairs and binaries (for reviews, see Komossa 2006; Popović 2012; De Rosa et al. 2019; Bogdanović et al. 2022). Dual AGN, i.e., galaxies with two unbound, actively accreting SMBHs, have been identified at \gtrsim kiloparsec separations (e.g., Koss et al. 2012; Chen et al. 2022, and references therein). However, spatially resolving the two SMBHs becomes increasingly challenging as their separation decreases. Spectroscopic features, such as the kinematic offset of AGN narrow lines, can also be used to identify AGN in merging galaxies (e.g., Comerford et al. 2009; Comerford & Greene 2014). To date, only one parsec-scale pair has been confirmed with very long baseline interferometry (Rodríguez et al. 2006; Bansal et al. 2017) despite large-scale searches (Burke-Spolaor 2011; Breiding et al. 2021).

Electromagnetic searches for subparsec SMBH binaries typically focus on features that encode the binary’s orbital motion on the temporal or spectral variability of AGN.

Searches for offset broad emission lines have been used to identify several hundred candidates (Tsalmantza et al. 2011; Eracleous et al. 2012; Ju et al. 2013; Shen et al. 2013), but this method is subject to false positives and other limitations (Gezari et al. 2007; Runnoe et al. 2015, 2017; Pflueger et al. 2018; Kelley 2021a). Periodically variable light curves (Farris et al. 2014; D’Orazio et al. 2015; Bowen et al. 2018; D’Orazio & Di Stefano 2018) have yielded a similar number of candidates (Graham et al. 2015; Charisi et al. 2016; Liu et al. 2019), though these samples are likely to also suffer significant contamination (Vaughan et al. 2016; Charisi et al. 2018; Sesana et al. 2018; Kelley et al. 2019b; Xin et al. 2020). Despite these challenges, the advent of large time domain surveys with the Vera Rubin observatory (Ivezic et al. 2019), combined with multiwavelength observations and increasing PTA sensitivity to the GWB, offers exciting opportunities for deriving multiple independent constraints on SMBH populations (e.g., Kelley et al. 2019a; Bogdanović et al. 2022). The prospects for low-frequency multimessenger astrophysics are discussed further in Section 5.

1.4. The Astrophysical Imprint on the GWB

All of the binary inspiral processes discussed above are imprinted on the GWB created by a population of SMBH binaries. Therefore, studying the GWB constitutes an important channel to obtain significant and novel insights on galaxy and binary mergers. For example, interactions with the binary environment and orbital eccentricities impact the shape of the GWB spectrum (Sesana 2013a). Stellar- and gas-driven binary hardening will cause a flattening or turnover of the low-frequency GWB spectrum relative to the single power law predicted for GW-only evolution (Kocsis & Sesana 2011). The primary effect of eccentricity is to boost GW emission to higher frequencies, owing to the emission of GWs at higher harmonics beyond twice the binary orbital frequency, which dominates for circular orbits (Enoki & Nagashima 2007). However, at extreme eccentricities ($\gtrsim 0.9$ – 0.95), close pericentric passages drive very rapid binary inspiral, leading to an overall attenuation of GWB amplitude at all frequencies (e.g., Kelley et al. 2017b).

The GW observations will also probe the history of SMBH mass growth. The GWB depends strongly on the distribution of binary chirp masses, \mathcal{M} , given by

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{M^{1/5}} = M \frac{q^{3/5}}{(1+q)^{6/5}}, \quad (1)$$

where $q \equiv m_2/m_1 < 1$ is the binary mass ratio, $M = m_1 + m_2$ is the total binary mass, and m_1 and m_2 are the masses of each SMBH. As a result, the GWB is intimately related to the SMBH mass function through its dependence on the chirp mass and, in turn, to the scaling relations of SMBH mass with host galaxy properties. These relations are well studied in the local Universe but unconstrained at higher redshifts; thus, detailed studies of the GWB will provide a novel path to probing these relations.

Previous stochastic GWB constraints have been used to probe the SMBH binary population by comparing with theoretical predictions for SMBH binary formation and evolution. All GWB results were strictly upper limits until the NANOGrav 12.5 yr data set, but the limits were still potentially constraining. The constraints were especially

informative when combined with electromagnetic observations of binary AGN candidates (Holgado et al. 2018; Inayoshi et al. 2018; Sesana et al. 2018; Nguyen et al. 2020). After the PPTA upper limit at 2.8 nHz (Shannon et al. 2013), it was first suggested that this ruled out a large range of SMBH binary model space (Shannon et al. 2015). However, Middleton et al. (2016) showed that the upper limits were consistent with a wide variety of plausible astrophysical models and that, in general, upper limits alone would be relatively unconstraining until they were about an order of magnitude smaller. Subsequent work showed the importance of analyzing many pulsars and accounting for their red noise and systematic errors in solar system ephemerides when establishing PTA upper limits (Arzoumanian et al. 2018; Hazboun et al. 2020; Vallisneri et al. 2020; Johnson et al. 2022). Since the 12.5 yr NANOGrav data set showed evidence for a common red-noise process consistent with (but not unambiguously attributable to) GWs, the measurement was shown to be consistent with a population of SMBH binaries with reasonable properties (Middleton et al. 2021).

1.5. Astrophysical Modeling of SMBH Binary GWB

Over the last few decades, a variety of different approaches have been used to model populations of SMBH binaries,⁷⁶ with a wide range of predictions for the resulting GWB amplitude. (See Appendix A for a summary of these model predictions and a comparison with the NANOGrav 15 yr results.) Many of these studies start from either semianalytic galaxy evolution models to obtain galaxy merger rates (Rajagopal & Romani 1995) or halo merger trees with added galaxies (Menou et al. 2001; Sesana et al. 2004) onto which an SMBH binary population model can be imposed. In lieu of physically modeling environmentally driven SMBH binary evolution, galaxy mergers are often directly linked to the formation of a close SMBH binary emitting GWs at PTA frequencies, and a power-law form is assumed for the GWB (e.g., Phinney 2001; Jaffe & Backer 2003; Wyithe & Loeb 2003; Enoki et al. 2004; Simon & Burke-Spolaor 2016). Some semianalytic models also include prescriptions for physical processes that cause GWB spectra to deviate from a pure power law, such as interactions of the binary with the gaseous and stellar environment of its host galaxy, discreteness of the binary population, and orbital eccentricity (e.g., Sesana et al. 2008, 2009; Sesana 2013b; McWilliams et al. 2014; Ravi et al. 2014; Bonetti et al. 2018b; Ryu et al. 2018; Chen et al. 2020). Versions of the semianalytic model approach have also been applied to catalogs of specific galaxies or quasars from observations (Rosado & Sesana 2014; Simon et al. 2014; Mingarelli et al. 2017; Casey-Clyde et al. 2022).

An alternative to the semianalytic modeling approach is to directly trace galaxy and SMBH evolution in cosmological hydrodynamics simulations (e.g., Kulier et al. 2015; Salcido et al. 2016; Kelley et al. 2017a, 2017b, 2018; Siwek et al. 2020; Volonteri et al. 2020; Curylo & Bulik 2022). This approach has the advantage of providing detailed information about the internal structures of galaxies and how they interact with SMBHs via AGN fueling and feedback. However, cosmological hydrodynamical simulations are very computationally

⁷⁶ Throughout this work, we will often use the term “SMBH binaries” to encompass SMBH pairs, even when the two SMBHs are not yet gravitationally bound but merely reside in the same galaxy.

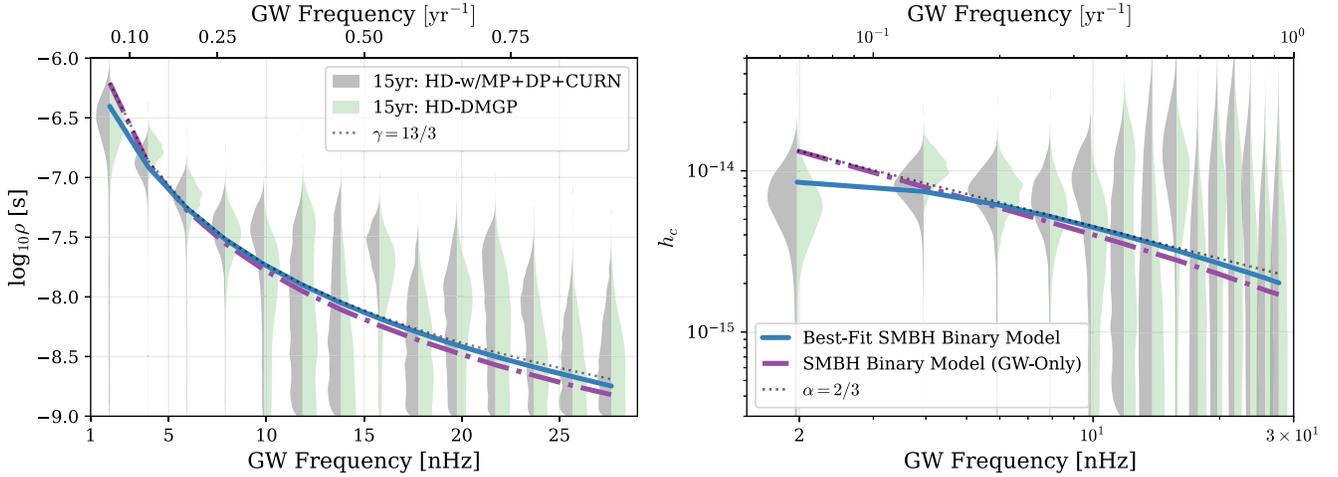


Figure 1. NANOGrav 15 yr GWB free-spectrum posteriors translated into the square root of the timing-residual power (ρ ; left panel) and characteristic strain (h_c ; right panel). The HD-correlated free spectrum measured while simultaneously fitting for MP, DP, and CURN free spectra (*HD-w/MP+DP+CURN*; gray violins, left side) is compared against the *HD-DMGP* model in which DM variations are modeled using GPs (green violins, right side). The black dotted lines show the idealized power-law spectra ($\rho^2 \propto f^{-13/3}$ and $h_c \propto f^{-2/3}$) fit to the median posterior value for the amplitude obtained from the *HD-w/MP+DP+CURN* model in [NG15gwb](#). Overplotted are the best-fitting simulated GWB spectra from models of SMBH binary populations produced in this analysis. Two models are shown, one that includes environmentally driven binary evolution (blue) and another that assumes GW-only evolution (purple). Both models are able to reproduce the data, while the environmentally driven model produces a slightly better fit. We conclude that the observed GWB spectrum is consistent with astrophysically motivated expectations from populations of SMBH binaries.

expensive compared to semianalytic models, and even the highest-resolution simulations must rely on subgrid prescriptions to model unresolved processes, including SMBH accretion, mergers, and feedback. Each of these complementary approaches therefore offers benefits and drawbacks, and importantly, each introduces certain systematics in their predictions for binary populations. In this work, we adopt a semianalytic modeling approach to SMBH binary population synthesis and defer the use of cosmological hydrodynamics simulations for future work.

1.6 Summary and Outline

Figure 1 shows the GWB spectrum recovered from the 15 yr NANOGrav data, along with the best-fitting simulated GWB spectra produced in this work. In Section 2, we summarize the NANOGrav 15 yr data set that forms the observational basis for this analysis and the GWB spectra derived from it (gray and green “violins”). In Section 3, we describe our methods of modeling populations of SMBH binaries and calculating the GWB spectra that they would produce. There, we also detail the approach that we use to compare our simulations to the 15 yr data. Our best-fitting models (colored curves) are presented in Section 4.

We find that astrophysically motivated models of SMBH binary populations are able to accurately reproduce the observed GWB spectrum (Section 4.1 and 4.2). We focus our analysis on two population models. One includes a self-consistent prescription for environmentally driven binary evolution (blue), and the other assumes GW-only evolution (purple), which is still commonly used in the literature. Both models are able to fit the data, while the environmentally driven case produces a slightly better match, particularly to the lowest frequency bin. We present the binary evolution parameters favored by 15 yr spectral fits for both models (Section 4.3). While the posterior distributions are broadly consistent with astrophysical expectations, the parameters tend to be shifted toward values that produce larger GWB amplitudes than was

previously most favored. Generally higher binary masses or densities or highly efficient binary mergers are required to produce the observed amplitudes. The characteristics of the implied binary populations are presented in Section 4.4.

Our results are discussed in the context of the field in Section 5, along with highlights for the near future of low-frequency GW astronomy.

Throughout this paper, we assume a WMAP9 cosmology with $\Omega_m = 0.228$, $\Omega_b = 0.0472$, and $H_0 = 0.6933 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

2. PTA Data

This work is based on the NANOGrav 15 yr data set, which includes 68 pulsars, 67 of which have a baseline of at least 3 yr and are included in the GWB analysis. The complete description of the data set can be found in Agazie et al. (2023b, hereafter [NG15](#)), while the detector characterization and noise modeling of individual pulsars is described in Agazie et al. (2023c, hereafter [NG15detchar](#)). The detailed description of the Bayesian search for the GWB is presented in [NG15gwb](#). Here we briefly summarize the measurement of the GWB spectrum from the NANOGrav data, focusing on the pieces that are necessary for the astrophysical interpretation presented in this paper.

The PTA collaborations systematically monitor millisecond pulsars and record the TOAs of their radio pulses. For each pulsar, a timing model is constructed, which estimates various factors affecting the TOA, including its astrometry (sky position, proper motion, and parallax), its spin period and spin period derivative, and binary parameters for pulsars with companions. Additionally, variations in the ionized interstellar medium along the line of sight, also known as the dispersion measure (DM), are included in our model. The analysis of each pulsar provides a best-fit estimate for the timing residuals, $r(t)$, which are the differences between the TOAs and the timing model. For more on the construction of the timing residuals in NANOGrav’s data set, see Section 4 of [NG15](#).

All red-noise processes, including the GWB itself, are modeled with a Fourier basis computed on the TOAs, as discussed in Section 2 of NG15gwb. The frequencies are $f_i = i/T_{\text{obs}}$, where $T_{\text{obs}} = 16.03$ yr is the time between the first and last TOA included in this data set.⁷⁷ The search for a GW signal is performed by constructing the cross-correlations of the residuals between pairs of pulsars, a and b , i.e.,

$$\langle r_a(t)r_b(t) \rangle \propto \int S_{ab}(f)df, \quad (2)$$

where f is the observer-frame GW frequency, and S_{ab} is the timing-residual cross-correlated power spectral density (PSD),

$$S_{ab}(f) = \Gamma(\xi_{ab})\Phi(f). \quad (3)$$

Here $\Phi(f)$ is the PSD of the timing residuals describing the spectrum of the process that is common among all pulsars, and Γ is the overlap reduction function, which describes the induced correlation between a pair of pulsars as a function of their angular separation, ξ_{ab} . The timing-residual PSD is related to the characteristic GW strain, $h_c(f)$, by

$$\Phi(f) = \frac{h_c(f)^2}{12\pi^2 f^3}. \quad (4)$$

The overlap reduction function is given by $\Gamma(\xi_{ab}) = \delta_{ab}$ for a CURN model and the characteristic HD pattern in the case of an isotropic GWB (Hellings & Downs 1983).

A specific spectral shape is typically prescribed to the common red-noise process (see Section 3.1). Traditionally, a power law has been used, and the detailed spectral analysis presented in NG15gwb shows support for this idealized, simple model. However, deviations appear at a variety of frequencies, which may skew the determination of a spectral slope (see Figures 1 and 6 in NG15gwb). It is therefore important to model the individual Fourier coefficients independently rather than enforcing a specific spectral shape on the PSD. The resulting “free spectrum” provides a minimally modeled Bayesian spectral characterization of PTA data. The free spectrum recovers the posterior of the common red-noise power spectrum at all sampling frequencies and is parameterized by the coefficient ρ , where $\rho_i^2 = \Phi(f_i)/T_{\text{obs}}$ is the power in the cross-correlated timing residuals.

Figure 1 shows the free-spectrum posteriors both in terms of the square root timing-residual power (ρ) and converted into GW characteristic strain (h_c). At the current signal-to-noise ratio, the spectral characterization of the signal is uncertain, and the recovered HD-correlated GWB signal may be impacted by non-HD-correlated noise. For our astrophysical interpretation, we adopt the free-spectrum posteriors from the 15 yr HD-correlated free spectrum modeled simultaneously with additional monopole-correlated (MP) and dipole-correlated (DP) red noise and CURN. This model, which we refer to as *HD-w/MP+DP+CURN* (gray), provides the most conservative constraints on the recovered GWB spectrum. As an additional comparison, we also analyze the HD-correlated free-spectrum posteriors utilizing an alternate model for DM variations, which we refer to as *HD-DMGP* (green), described in detail in Section 5.1 of NG15gwb.

Figure 1 also shows the median posterior amplitude value for the idealized power-law fit to the *HD-w/MP+DP+CURN* model. Even though the power-law model provides an illustrative example for quick model comparison, we do not include it in the astrophysical interpretation because it fails to encapsulate the full range of information contained in the free-spectrum posteriors.

The number of Fourier components used in an analysis is typically chosen based on the preference of the data for various red-noise processes (e.g., in Arzoumanian et al. 2020, the CURN model preferred only five frequencies, while in NG15gwb, that number increased to 14). While the CURN model prefers 14 Fourier components in the 15 yr data set, the HD-correlated free-spectrum posteriors provide strong constraints only in the five lowest frequency bins; thus, only those bins are used in this analysis. However, we find no difference in our results if we expand to using the full 14 frequencies.

3. Methods

Our goal is to constrain the properties of the underlying SMBH binary population that can produce a GWB consistent with the NANOGrav 15 yr data. Our approach consists of three main components described below and depicted schematically in Figure 2.

SMBH binary population synthesis simulations (Section 3.1–3.3)—We generate “libraries” of SMBH binary populations and their GW signals, exploring a large range of the binary formation/evolution parameter space. For this, NANOGrav has developed a flexible framework for SMBH binary population synthesis called *holodeck* (L. Z. Kelley et al. 2023, in preparation),⁷⁸ which allows us to explore the binary population models and encompass systematic uncertainties. Within *holodeck*, we determine the number density of the cosmic population of SMBH binaries using semianalytical models based on the observationally constrained properties of galaxies and galaxy mergers. Using an SMBH–host relation, specifically the correlation between the mass of the SMBH and the mass of the stellar bulge, i.e., $M_{\text{BH}}-M_{\text{bulge}}$, we assign SMBH masses to the mergers and calculate the binary evolution from large separations down to the GW regime. From each population, we compute the GWB signals they would produce.

Interpolation of the population synthesis models (Section 3.4)—The simulated GWB spectra are sampled at discrete points of the multidimensional binary population parameter spaces that we explore. We refer to the collection of simulated spectra for a given parameter space as a “library.” We then use GPs to interpolate between the population synthesis simulations and predict the shape of the GWB spectrum for any point of the parameter domain. This is necessary because the population simulations are too computationally expensive to run live while fitting against the NANOGrav data.

Fitting population synthesis models against PTA data (Section 3.5)—We use a Markov Chain Monte Carlo (MCMC) approach to fit the trained GPs against the input free-spectrum posteriors from NG15gwb, generating posterior distributions of the binary population model parameters. From these, we constrain the different SMBH binary populations and evolutionary scenarios that could produce the observed GWB.

⁷⁷ This data set is named “15 yr data set” because no single pulsar exceeds 16 yr of observations, even though the total time spanned by the entire set of observations is 16.03 yr.

⁷⁸ <https://github.com/nanograv/holodeck>

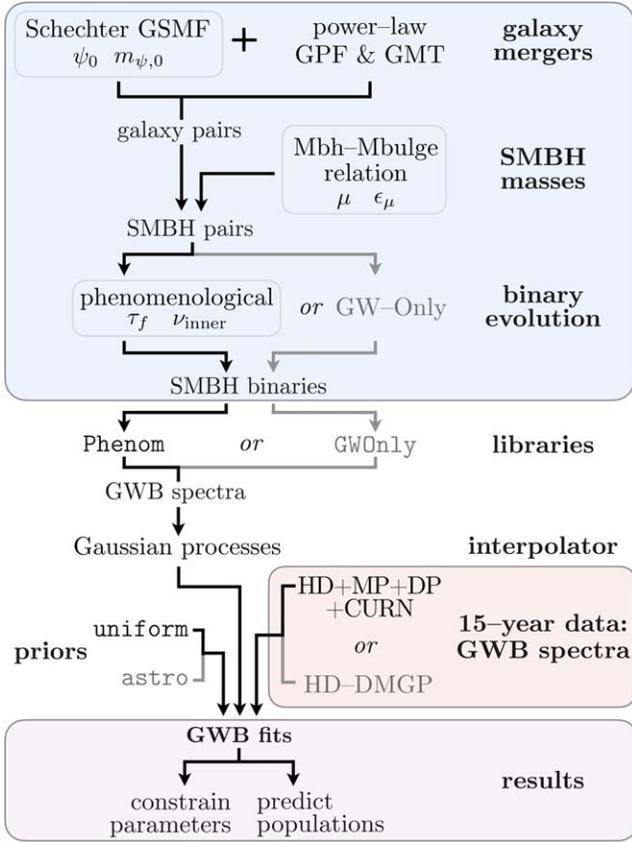


Figure 2. Schematic of the analysis pipeline used in this study. Our population models are composed of galaxy mergers combined with SMBH masses and a prescription for binary evolution. Galaxy merger rates are derived from the combination of a GSMF, GPF, and GMT. Binary evolution can either follow the self-consistent, phenomenological approach or assume GW-only evolution. Libraries contain a large number of binary populations and their resulting GWB signals, which are calculated for varying uncertain physical parameters in both evolution scenarios. Gaussian processes (GPs) are used to interpolate across the library parameters when fitting against the NANOGrav 15 yr data. Fits can be performed using broad uniform priors or with more tightly constrained, astrophysically motivated priors. Fits are performed against the 15 yr data using GWB spectra derived from the *HD-w/MP+DP+CURN* or *HD-DMGP* models (see text).

3.1. GWs from SMBH Binary Populations

The GWB spectrum can be calculated as the integrated GW emission of individual binaries throughout the Universe. The characteristic strain of the GWB over a given logarithmic interval of frequency can be expressed as (Phinney 2001; Wyithe & Loeb 2003)

$$h_c^2(f) = \int dM dq dz \frac{\partial^4 N}{\partial M \partial q \partial z \partial \ln f_p} h_s^2(f_p). \quad (5)$$

The sky- and polarization-averaged GW spectral strain from a single circular binary h_s can be related to a binary's total GW luminosity, L_{GW} , as (Finn & Thorne 2000)

$$h_s^2(f) = \frac{G}{c^3} \frac{L_{\text{GW}}}{(2\pi f_p)^2 d_c^2} = \frac{32}{5c^8} \frac{(GM)^{10/3}}{d_c^2} (2\pi f_p)^{4/3}. \quad (6)$$

Here d_c is the comoving distance to a source at redshift z . Because the GW frequency is twice the orbital frequency for circular binaries, the observer-frame GW frequency f can be related to the rest-frame orbital frequency f_p as $f = 2f_p/(1+z)$.

Throughout this paper, we take the chirp mass \mathcal{M} (Equation (1)) and, by extension, the total binary mass M to be intrinsic rest-frame properties of the binary.

In practice, it is much more convenient to calculate a comoving volumetric number density of binaries $\eta \equiv dN/dV_c$ and use this quantity to infer the full population (Rajagopal & Romani 1995; Jaffe & Backer 2003; Sesana et al. 2008):

$$\begin{aligned} \frac{\partial^4 N}{\partial M \partial q \partial z \partial \ln f_p} &= \frac{\partial^3 \eta}{\partial M \partial q \partial z} \frac{\partial t}{\partial \ln f_p} \frac{\partial z}{\partial t} \frac{\partial V_c}{\partial z}, \\ &= \frac{\partial^3 \eta}{\partial M \partial q \partial z} \cdot \tau(f_p) \cdot 4\pi c (1+z) d_c^2. \end{aligned} \quad (7)$$

Here $\tau(f_p) \equiv f_p/(df_p/dt)$ is the binary hardening timescale, the rest-frame duration that a binary spends in a given logarithmic interval of frequency. Equation (7) connects the redshift evolution to the time evolution of binary sources over frequencies. For a circular binary evolving purely due to GW emission, the rate of semimajor axis change and the hardening timescale are given by Peters (1964):

$$\begin{aligned} \left. \frac{\partial a}{\partial t} \right|_{\text{gw}} &= -\frac{64 G^3 m_1 m_2 M}{5 c^5 a^3}, \\ \tau_{\text{gw}} \equiv \frac{\partial t}{\partial \ln a} &= \frac{5}{96} \left(\frac{GM}{c^3} \right)^{-5/3} (2\pi f_p)^{-8/3}. \end{aligned} \quad (8)$$

Combining the above equations with the comoving volume of a light cone (e.g., Hogg 1999),

$$\frac{\partial V_c}{\partial z} \frac{\partial z}{\partial t} = 4\pi c (1+z) d_c^2, \quad (9)$$

gives the idealized expression for a GWB produced by circular, GW-only driven SMBH binaries (Phinney 2001):

$$\begin{aligned} h_{s,\text{circ}}^2(f) &= \frac{4\pi}{3c^2} (2\pi f)^{-4/3} \\ &\int dM dq dz \frac{\partial^3 \eta}{\partial M \partial q \partial z} \frac{(GM)^{5/3}}{(1+z)^{1/3}}. \end{aligned} \quad (10)$$

This motivates the common expression for GWB spectra as a power law of the form

$$h_c(f) = A_{\text{yr}} \cdot (f/\text{yr}^{-1})^{-\alpha}, \quad (11)$$

where A_{yr} is the GWB amplitude referenced at a frequency of 1 yr^{-1} , and in the idealized case, $\alpha = 2/3$. Because the timing-residual PSD of a GW signal is related to the characteristic GW strain by Equation (4), this ideal power-law form of the GWB can be expressed relative to a reference frequency f_{ref} equivalently as

$$\Phi(f) = \frac{A^2}{12\pi^2} \left(\frac{f}{f_{\text{ref}}} \right)^{-\gamma} f_{\text{ref}}^{-3}. \quad (12)$$

Note that we have defined the power-law indices to be positive quantities such that $h_c \propto f^{-\alpha}$ and $\Phi \propto f^{-\gamma}$. The power-law indices are therefore related as $\gamma = 3 + 2\alpha$, such that the idealized, GW-only index is $\gamma = 13/3$.

Realistic GWB spectra can deviate substantially from a power law, primarily due to the following three effects.

Interactions with the binary environment—Astrophysical processes that extract energy and angular momentum from the

binary (e.g., via stellar and gaseous interactions) can accelerate its frequency evolution relative to the GW-only evolution. Therefore, any binary hardening via processes other than GW emission will necessarily result in an attenuation of the GWB compared to a purely GW-driven spectrum, as binaries spend less time emitting GWs in a given frequency interval. This effect is particularly important at low frequencies ($f \ll 1 \text{ yr}^{-1}$), where binaries can more easily couple to their local galactic environments (Begelman et al. 1980; Kocsis & Sesana 2011) and GW emission is weaker. In fact, coupling between SMBHs and their astrophysical environments is required for binaries to reach the PTA band within a Hubble time. The question is thus whether the resulting flattening (or turnover) in the GWB spectrum occurs within the PTA band or at frequencies too low to be currently accessible.

Discreteness of the binary population—Equation (7) assumes a continuous distribution of SMBH binaries across the (M, q, z, f) parameter space. At low frequencies ($f \lesssim 1 \text{ yr}^{-1}$), the hardening timescale is very long, and a large number of binaries contribute to the GWB, making this approximation valid. At higher frequencies ($f \gtrsim 1 \text{ yr}^{-1}$), however, the hardening timescale becomes shorter, and the typical number of binaries producing the bulk of the GWB energy in a given frequency bin approaches unity (Sesana et al. 2008). In this regime, a continuous distribution overestimates the GWB signal. Properly accounting for the finite number of sources in each frequency bin therefore results in a steeper GWB spectrum at high frequencies (Sesana et al. 2009). While a given overall amplitude of the GWB can be produced by either a larger number of lower-mass SMBH binaries or a smaller number of higher-mass binaries, these differences change the frequency at which discreteness becomes important. As a result, they change the location and severity of the high-frequency spectral steepening.

Orbital eccentricity—Unlike circular binaries that emit GWs at exactly twice the orbital frequency, eccentric binaries emit GW energy at all integer harmonics. This leads to GW energy being moved from lower to higher frequencies (Enoki et al. 2004). Additionally, smaller pericenter distances tend to increase the rate of binary inspiral. These factors produce a variety of effects, including a spectral turnover at low frequencies, a flatter spectrum at higher frequencies, and a “bump” in between (Enoki & Nagashima 2007; Sesana 2013a; Huerta et al. 2015; Chen et al. 2017; Kelley et al. 2017b). However, for these effects to be substantial, very large eccentricities ($e \gtrsim 0.9$) are necessary at very small separations (well within the PTA band).⁷⁹ Since this is not expected to be the case, we restrict the current analysis to circular binaries.

These effects highlight the additional information encoded in the deviations of the GWB spectra from a pure power law and the importance of careful modeling of the binary population. The above considerations also demonstrate the need for explicit integration of the binary evolution that includes environmental interactions, the discreteness of binaries, and their expected cosmic variance.

⁷⁹ While recent results suggest that circumbinary accretion disks may drive moderate eccentricities ($e \sim 0.4\text{--}0.5$) in some systems (D’Orazio & Duffell 2021; Zrake et al. 2021; Siwek et al. 2023), the effects are unlikely to be detectable in the GWB. Such processes could be more important for individually detectable GW signals, particularly rapidly accreting ones that may be promising multimessenger sources.

3.2. SMBH Binary Population Synthesis

Many previous works have constructed model populations of SMBH binaries and obtained predictions for the resulting GWB. These have produced a wide range of predictions for the GWB amplitudes, which are summarized in Table A1. These SMBH binary population models generally involve either semianalytic models or cosmological hydrodynamic simulations. In this work, we focus only on semianalytic models and defer the exploration of binary populations from cosmological simulations to a future study.

In general, three key components are required for modeling the binary populations responsible for the GWB: (i) galaxy masses and merger rates (Section 3.2.1), (ii) SMBH masses based on a galaxy–host relationship (Section 3.2.2), and (iii) a binary evolution prescription (Section 3.2.3). We choose particular parameterizations for each of these components following Chen et al. (2019), described below, which are implemented in the `holodeck` code. A large number of free parameters are required for any such type of population synthesis calculation, more than can be meaningfully fit by the existing data. We therefore identify key parameters to vary in the models considered here and adopt standard literature values for the rest. In our analysis, we use these models to construct numerous different libraries of binary populations, and we explore the impact of varying these key parameters on the resulting GWB spectra.

For the binary evolution, we consider libraries using both a phenomenological binary inspiral model (dubbed *Phenom*) and a naive GW-only inspiral scenario (*GWOnly*). In both cases, the libraries vary two parameters that determine the galaxy number density (the normalization ψ_0 and turnover mass $m_{\psi,0}$) and two parameters describing the $M_{\text{BH}}\text{--}M_{\text{bulge}}$ relationship (the normalization μ and intrinsic scatter ϵ_μ). The *Phenom* library includes two additional parameters describing the total binary lifetimes τ_f and the binary hardening rate at small separations ν_{inner} . These models and parameters are described in detail in the following sections. Table B1 lists all of the parameters, giving fiducial values when they are fixed and the prior distributions for those that are varied. Table B2 summarizes our different libraries and which parameters are varied in each. In Appendix C, we also compare against larger, “extended” models in which additional parameters are varied (*Phenom-Ext* and *GWOnly-Ext*).

3.2.1. The Galaxy Merger Rate

The number density of galaxy mergers ($\eta_{\text{gal-gal}}$) can be expressed (Chen et al. 2019) in terms of a galaxy stellar-mass function (GSMF; Ψ), galaxy pair fraction (GPF; P), and galaxy merger time (GMT; $T_{\text{gal-gal}}$):

$$\frac{\partial^3 \eta_{\text{gal-gal}}}{\partial m_{*1} \partial q_* \partial z} = \frac{\Psi(m_{*1}, z')}{m_{*1} \ln(10)} \frac{P(m_{*1}, q_*, z')}{T_{\text{gal-gal}}(m_{*1}, q_*, z')} \frac{\partial t}{\partial z'}. \quad (13)$$

This distribution is calculated in terms of the stellar mass of the primary galaxy m_{*1} , the stellar mass ratio ($q_* = m_{*2}/m_{*1} \leq 1$), and the redshift z . Because the galaxy merger spans a finite timescale ($T_{\text{gal-gal}}$) and corresponding redshift interval, we distinguish between the initial redshift at which a galaxy pair forms ($z' = z'[t]$ at some initial time t) and the redshift at which the system becomes a postmerger galaxy remnant ($z = z[t + T_{\text{gal-gal}}]$). An additional delay is required for binaries

to reach the PTA frequency band, which is characterized in Section 3.2.3.

The GSMF is defined as

$$\Psi(m_{\star 1}, z') \equiv \frac{\partial \eta_{\star}(m_{\star 1}, z')}{\partial \log_{10} m_{\star 1}}, \quad (14)$$

i.e., the differential number density of galaxies per decade of stellar mass. The implementation used in this analysis described the GSMF in terms of a single Schechter function (Schechter 1976),

$$\Psi(m_{\star 1}, z) = \ln(10) \Psi_0 \cdot \left[\frac{m_{\star 1}}{M_{\psi}} \right]^{\alpha_{\psi}} \exp\left(-\frac{m_{\star 1}}{M_{\psi}}\right), \quad (15)$$

where we have introduced Ψ_0 , M_{ψ} , and α_{ψ} as new variables. In order to allow the GSMF to vary with redshift, we parameterize these quantities as

$$\begin{aligned} \log_{10}(\Psi_0/\text{Mpc}^{-3}) &= \psi_0 + \psi_z \cdot z, \\ \log_{10}(M_{\psi}/M_{\odot}) &= m_{\psi,0} + m_{\psi,z} \cdot z, \\ \alpha_{\psi} &= 1 + \alpha_{\psi,0} + \alpha_{\psi,z} \cdot z, \end{aligned} \quad (16)$$

such that each of these quantities has a simple linear scaling with redshift. This introduces six new dimensionless parameters into our models, corresponding to the normalization (ψ_0 , $m_{\psi,0}$, and $\alpha_{\psi,0}$) and slope (ψ_z , $m_{\psi,z}$, and $\alpha_{\psi,z}$) of the redshift scaling. In all of the analyses presented here, the latter three are always kept fixed at the fiducial values specified in Table B1. The GSMF normalization and characteristic mass parameters ψ_0 and $m_{\psi,0}$ are allowed to vary in our fiducial *Phenom* library, while $\alpha_{\psi,0}$ is additionally varied in *Phenom-Ext*.

The GPF and GMT are defined as

$$P(m_{\star 1}, q_{\star}, z') \equiv \frac{\partial}{\partial q_{\star}} \left(\frac{N_{\star, \text{pairs}}(m_{\star 1}, q_{\star}, z')}{N_{\star}(m_{\star 1}, z')} \right), \quad (17)$$

$$T_{\text{gal-gal}}(m_{\star 1}, q_{\star}, z') \equiv \int_{a_{\star,i}}^{a_{\star,f}} [\dot{a}_{\star}(m_{\star 1}, q_{\star}, z')]^{-1} da_{\star}, \quad (18)$$

where \dot{a}_{\star} denotes the rate at which the merging galaxies' separation decreases. The GPF describes the number of observable galaxy pairs relative to the number of all galaxies. The GMT is the duration over which two galaxies can be discernible as pairs from an initial separation $a_{\star,i}$ at which they are associated with one another until a final separation $a_{\star,f}$, after which they are no longer distinguishable as separate galaxies. These two distributions are typically determined empirically based on the detection of galaxy pairs in observational surveys and thus depend on observational definitions and selection criteria (e.g., Conselice et al. 2008; Mundy et al. 2017; Snyder et al. 2017; Duncan et al. 2019).

In practice, we parameterize $P(m_{\star 1}, q_{\star}, z')$ and $T_{\text{gal-gal}}(m_{\star 1}, q_{\star}, z')$ as redshift-dependent power laws of $m_{\star 1}$, q_{\star} , and z following Chen et al. (2019):

$$\begin{aligned} P(m_{\star 1}, q_{\star}, z') &= P_0 \left(\frac{m_{\star 1}}{10^{11} M_{\odot}} \right)^{\alpha_p} (1+z)^{\beta_p} q^{\gamma_p} \\ \alpha_p &= \alpha_{p,0} + \alpha_{p,z} \cdot z \\ \gamma_p &= \gamma_{p,0} + \gamma_{p,z} \cdot z, \end{aligned} \quad (19)$$

$$\begin{aligned} T_{\text{gal-gal}}(m_{\star 1}, q_{\star}, z') &= T_0 \left(\frac{m_{\star 1}}{10^{11} M_{\odot}/h} \right)^{\alpha_t} (1+z)^{\beta_t} q^{\gamma_t} \\ \alpha_t &= \alpha_{t,0} + \alpha_{t,z} \cdot z \\ \gamma_t &= \gamma_{t,0} + \gamma_{t,z} \cdot z. \end{aligned} \quad (20)$$

As shown in Table B1, the values of the corresponding parameters are kept fixed to the standard literature values in our fiducial model (*Phenom*), but the parameters governing the scaling of the GPF and GMT with redshift ($\beta_{p,0}$ and $\beta_{t,0}$), the scaling of the GPF with mass ratio ($\gamma_{p,0}$), and the GMT normalization (T_0) are allowed to vary in our extended models, as described in more detail below.

3.2.2. The SMBH–Host Relation

The number density of galaxy mergers given in Equation (13) is a distribution that describes the number of galaxy pairs as a function of galaxy properties. We assume a one-to-one correspondence between galaxy pairs and SMBH binaries and adopt an SMBH–host relationship to translate from galaxies to SMBHs. In this analysis, we restrict ourselves to the $M_{\text{BH}}-M_{\text{bulge}}$ relationship, which relates the galaxy stellar bulge mass to the SMBH mass for each component of the binary as (Marconi & Hunt 2003)

$$\log_{10}(M_{\text{BH}}/M_{\odot}) = \mu + \alpha_{\mu} \log_{10} \left(\frac{M_{\text{bulge}}}{10^{11} M_{\odot}} \right) + \mathcal{N}(0, \epsilon_{\mu}). \quad (21)$$

Here $\mathcal{N}(0, \epsilon_{\mu})$ denotes the normally distributed random scatter with a mean of zero and standard deviation of ϵ_{μ} (in dex). This relation depends on three model parameters that are allowed to vary in our analyses: the dimensionless black hole mass normalization (μ), intrinsic scatter (ϵ_{μ} , in dex), and power-law index α_{μ} (which is varied only in our extended models and is dimensionless). A fraction of the galaxy stellar mass is in the stellar bulge component ($M_{\text{bulge}} = f_{\star, \text{bulge}} \cdot m_{\star}$), which we take to be $f_{\star, \text{bulge}} = 0.615$ based on empirical bulge fraction measurements of massive galaxies from Lang et al. (2014) and Bluck et al. (2014). Using the $M_{\text{BH}}-M_{\text{bulge}}$ relationship, we transform the number density of galaxy mergers to a number density of SMBH binaries via

$$\frac{\partial^3 \eta}{\partial M \partial q \partial z} = \frac{\partial^3 \eta_{\text{gal-gal}}}{\partial m_{\star 1} \partial q_{\star} \partial z} \frac{\partial m_{\star 1}}{\partial M} \frac{\partial q_{\star}}{\partial q}. \quad (22)$$

Equation (22) provides an expectation value for the number of binaries in a point (M, q, z) in parameter space. To discretize the SMBH binary population and measure the effects of cosmic variance, we assume that the true number of binaries in any given spatial volume is Poisson-distributed. We then integrate the differential number of binaries over finite bins of parameter space to obtain the expected number of binaries in each bin. We generate multiple realizations by drawing many times from a Poisson distribution (\mathcal{P}) centered at that value. Finally, we sum over parameter space bins to calculate the resulting GWB spectrum. In practice, we

implement Equation (5) as

$$h_c^2(f) = \sum_{M,q,z,f} \mathcal{P} \left(\frac{\partial^4 N}{\partial M \partial q \partial z \partial f} \Delta M \Delta q \Delta z \Delta \ln f \right) \frac{h_s^2(f_p)}{\Delta \ln f}. \quad (23)$$

3.2.3. Binary Evolution

The final component for constructing the GWB is the most uncertain: the binary evolution from the initial galaxy merger until the eventual SMBH coalescence. Typically, interactions with the astrophysical environment (i.e., stars and gas in the host galaxy) are required to bring an SMBH binary into the PTA band within a Hubble time. For example, the high-mass binaries ($\mathcal{M} \gtrsim 10^9 M_\odot$) that dominate the GW signals in the PTA band must reach separations of ~ 0.1 pc before GW emission becomes dominant and drives efficient inspiral. Those binaries enter the NANOGrav band (currently $\sim 1/15 \text{ yr}^{-1} \approx 2 \text{ nHz}$) when they reach separations of ≈ 0.05 pc—only a factor of 2 smaller. This immediately implies that the environmental processes may play a nonnegligible role in binary evolution, even after the binaries reach the NANOGrav band.

Even if environmental hardening is effective in bringing binaries to the PTA-detectable frequencies, binary lifetimes can still be many billions of years, and a large fraction of binaries may stall (Kelley et al. 2017a). This can lead to binaries reaching the PTA band at substantially lower redshifts than those at which their respective galaxy mergers occurred.

Detailed modeling of environmentally driven binary evolution can introduce dozens of free parameters, even when the SMBH and galaxy parameters are known a priori. Ultimately, many of these parameters become significantly degenerate in determining the resulting shape of the GWB spectrum and the properties of the SMBH binaries producing it. For this reason, we focus this analysis on a “phenomenological” model that is designed to capture the overall effects of more explicit binary evolution while introducing only a small number of free parameters. In these models, the hardening timescale is parameterized in terms of the evolution of the binary semimajor axis a as

$$\left. \frac{da}{dt} \right|_{\text{phenom}} = H_a \cdot \left(\frac{a}{a_c} \right)^{1-\nu_{\text{inner}}} \cdot \left(1 + \frac{a}{a_c} \right)^{\nu_{\text{inner}}-\nu_{\text{outer}}}. \quad (24)$$

The hardening timescale is thus a double power law, with a break at the critical separation a_c and asymptotic behaviors of

$$\frac{dt}{d \ln a} (a \ll a_c) \sim a^{\nu_{\text{inner}}} \quad (25)$$

in the “inner” (small-separation) regime and

$$\frac{dt}{d \ln a} (a \gg a_c) \sim a^{\nu_{\text{outer}}} \quad (26)$$

in the “outer” (large-separation) regime. Hardening rates are added linearly, such that the total rate of evolution when also including GW emission (Equation (8)) is given by $da/dt = [da/dt]_{\text{phenom}} + [da/dt]_{\text{GW}}$. We assume a fixed value of $\nu_{\text{outer}} = +2.5$ in all of our analyses, motivated by detailed literature models of the dynamical friction–driven evolution of SMBH binaries (Kelley et al. 2017a). In our models, ν_{inner} ,

which controls the hardening rate of binaries as they approach and enter the PTA band, is allowed to vary.

In addition to the two power-law indices (ν_{inner} , ν_{outer}) and the characteristic break separation (a_c), the normalization (H_a) is calculated such that the total lifetime of the binary matches a target τ_f , i.e.,

$$\tau_f = \int_{a_{\text{init}}}^{a_{\text{isco}}} \left(\frac{da}{dt} \right)^{-1} da, \quad (27)$$

where a_{init} is the initial binary separation, and $a_{\text{isco}} \equiv 6 GM/c^2$ is the innermost stable circular orbit, where we consider the two SMBHs to have merged. While this expression for a_{isco} is based on the test-particle approximation ($q \ll 1$), the true value should differ by less than a factor of 2 (Flanagan & Hughes 1998) for low SMBH spins, and the contribution to the total lifetime is always negligible for $a \sim a_{\text{isco}}$. The total lifetime τ_f is a key parameter that we vary in our models.

At numerous points in our analysis, we compare the self-consistent phenomenological model (in the *Phenom* and *Phenom-Ext* libraries) against a model where binaries decay only due to GW emission (*GWOnly* and *GWOnly-Ext* libraries). In the GW-only model, we take the redshift (and thus source distance) to be the post-galaxy-merger redshift without an additional delay and set the binary evolution time in Equation (7) to be that of GW-only evolution (i.e., Equation (8)). This model is not self-consistent, as GW-only evolution is unable to bring binaries to the PTA band within a Hubble time. It is nonetheless a useful comparison because the GW-only assumption is often still used in the literature and tends to produce the highest GWB amplitudes.

Figure 3 shows the binary evolution and GWB spectra resulting from the phenomenological evolution model. Total binary lifetimes of 0.1 and 1 Gyr are plotted with solid and dashed lines, respectively, while varying small-separation power-law indices (ν_{inner}) are shown with different colors. In each panel, the medians and 50% interquartile ranges of the binaries are shown. Note that in the top panel, only binaries with $3 \times 10^8 M_\odot < M \leq 3 \times 10^9 M_\odot$ and $0.1 < q \leq 1.0$ are shown. In the environmentally driven regime (larger separations), their hardening rate is determined such that their total lifetime matches the target value. The narrow interquartile regions in the environmental regime reflect the small variations in hardening rate required to produce the target total lifetime for this range of masses.

The GW hardening rate, which dominates at small separations, is determined entirely by the binary masses for our assumption of circular orbits. In the phenomenological model, the hardening rate at larger separations is determined such that the total inspiral time matches the input binary lifetime. This means that shorter-lifetime populations are forced to transition into the GW-driven regime at smaller separations. The power-law indices also affect the transition point by determining which separations dominate the binary’s evolution time. More positive values of ν_{inner} lead to flatter evolution trajectories with less and less time spent at subparsec binary separations. For reference, the dotted vertical line shows the separation at which a binary with $M = 10^9 M_\odot$ enters the 15 yr NANOGrav frequency band. The two models with $\nu_{\text{inner}} = -0.5$ and -1.0 lead to environmentally driven evolution until sufficiently small separations so that the resulting GWB spectral turnover (bottom panel) is clearly visible in simulated 15 yr spectra.

We note that a value of $\nu_{\text{inner}} = -1.0$ is well motivated by numerical stellar scattering experiments of closely bound SMBH binaries (Sesana 2010; Sesana & Khan 2015). However, the true rate of environmental hardening for close binaries will depend on the stellar distribution in a given host galaxy, as well as the role of gas-driven binary evolution, motivating the choice to allow ν_{inner} to vary in our models.

The $\nu_{\text{inner}} = -0.5$ variation produces a substantial attenuation of the GWB: up to a 50% decrease in characteristic strain (75% reduction in GW power). Even though the 1.0 Gyr lifetime model (dashed lines) qualifies as efficient and rapid binary evolution, the overall amplitude of the GWB at all frequencies is $\sim 10\%$ lower than the GW-only model. This is because a fraction of the binaries, specifically those that formed within a look-back time of 1.0 Gyr, are unable to reach the PTA band before redshift zero and thus do not contribute to the observable GWB. This figure highlights that only in a narrow region of parameter space do realistic GWB spectra match predictions from GW-only models, but the differences between the self-consistent and GW-only models can be subtle.

3.3. Libraries of SMBH Binary Populations and GWB Spectra

With the models described above, we use the `holodeck` code to calculate libraries of SMBH binary populations and their resulting GWB spectra. In each parameter space that we explore, we include a large number of sample points in the space in addition to many realizations of populations and spectra at each point. We use the same GW frequency bins as the 15 yr NANOGrav data ($f_i = i/16.03 \text{ yr}^{-1} = i \times 1.98 \text{ nHz}$; see Section 2) to calculate spectra. Here we present the two primary libraries used in our analysis, *Phenom* and *GWOnly*, and outline some of their features. In Table B2, we summarize the parameters that are varied in each library, while the full list of model parameters (including fiducial values for fixed parameters and assumed prior distributions for varied parameters) is given in Table B1.

Tens of free parameters are required in these models, many of which are poorly constrained either observationally or theoretically. In addition, many of them are formally degenerate in their effects on the resulting GWB spectra (e.g., Chen et al. 2019). For this reason, we adopt as our fiducial library *Phenom*, a model with six parameters $\{\psi_0, m_{\psi,0}, \mu, \epsilon_\mu, \tau_f, \text{ and } \nu_{\text{inner}}\}$ that produce SMBH binary populations and GWB spectra that effectively span the broader model uncertainties. More specifically, this library varies the normalization and turnover mass of the GSMF (ψ_0 and $m_{\psi,0}$), along with the normalization and scatter of the $M_{\text{BH}}-M_{\text{bulge}}$ relationship (μ and ϵ_μ). The library also utilizes the phenomenological binary evolution model and varies the SMBH binary lifetime, τ_f , and the hardening power-law index at small separations ν_{inner} .

The differences in GWB spectra for systematic variations in the *Phenom* model parameters are shown in Figure 4. The overall amplitude of the GWB spectrum varies most significantly in the left and top middle panels, indicating that the GWB amplitude is most sensitive to the parameters determining SMBH masses ($m_{\psi,0}$, μ) and the SMBH binary number density (ψ_0). In the bottom middle panel, we see that increasing scatter in the $M_{\text{BH}}-M_{\text{bulge}}$ relationship (ϵ_μ) also increases the GWB amplitude. This is due to the fact that larger scatter increases the effective SMBH masses through Eddington bias; because low-mass SMBHs are more numerous, their scatter

toward higher masses outnumbers the scatter of the rarer, higher-mass SMBHs toward lower values. Notice that variations in SMBH mass, GSMF turnover mass, and $M_{\text{BH}}-M_{\text{bulge}}$ scatter (parameterized by μ , $m_{\psi,0}$, and ϵ_μ , respectively) all produce qualitatively similar changes in the GWB spectra. Higher masses preferentially increase the low-frequency amplitudes, thereby steepening the spectra at higher frequencies ($f \gtrsim 1 \text{ yr}^{-1}$). This occurs because rare high-mass binaries contribute less to the GWB at higher frequencies due to their rapid evolution at smaller separations (see Section 3.2). The ϵ_μ parameter shows this frequency-dependent effect even more prominently, as it preferentially affects the highest SMBH mass bins, where the gradient in SMBH number density with respect to mass is steepest.

The shape of the spectrum at low frequencies is determined by the binary hardening rate da/dt (as introduced in Section 3.2.3), which includes the interaction of binaries with their nuclear galactic environments. Recall that in our models, the binary lifetime is an input parameter, one that is varied in the top right panel of Figure 4 and kept fixed at $\tau_f = 1.0 \text{ Gyr}$ in all other panels. Consequently, for a given τ_f , binaries of different masses enter the GW regime at different frequencies. Some of the variations in low-frequency spectral shape seen when the mass-determining parameters (μ , ϵ_μ , and $m_{\psi,0}$) are varied and hardening parameters are kept fixed can therefore be attributed to τ_f being the same for all binary masses. For our models with fixed total binary lifetimes, populations with lower masses tend to have stronger low-frequency turnovers as lower-mass binaries enter the GW regime at higher frequencies. Equivalently, lower-mass systems spend more time at higher frequencies, meaning that their environmentally driven evolution must have proceeded even faster at lower frequencies.

While there is some degeneracy across all parameters, only the two parameters that directly affect the average binary mass ($m_{\psi,0}$ and μ) produce mostly degenerate spectral changes. As mentioned above, even the $M_{\text{BH}}-M_{\text{bulge}}$ scatter parameter (ϵ_μ), which also changes the average binary mass, is noticeably distinct. This speaks to the possibility of independently constraining multiple parameters with a sufficiently high signal-to-noise ratio even without appealing to additional information content such as sky anisotropy, individual continuous-wave sources, or electromagnetic counterparts and other multimessenger constraints.

As introduced in Section 3.2.3, in addition to *Phenom*, we use a library with the same variations in GSMF and $M_{\text{BH}}-M_{\text{bulge}}$ parameters but using GW-only evolution instead of the phenomenological model. We refer to this four-dimensional parameter space $\{\psi_0, m_{\psi,0}, \mu, \epsilon_\mu\}$ as *GWOnly*. Figure 5 shows a comparison of GWB spectra with variations in the *Phenom* model parameters versus spectra from the *GWOnly* models. For longer binary lifetimes ($\gg \text{Gyr}$), fewer systems are able to coalesce, and the GWB amplitude is noticeably diminished at all frequencies. For shorter lifetimes ($\lesssim \text{Gyr}$), non-GW hardening is still important at low frequencies within the 15 yr NANOGrav band. This leads to binaries evolving faster than the GW-only prediction, fewer binaries existing at these frequencies, and thus attenuated GW emission producing a low-frequency turnover (Kocsis & Sesana 2011; Ravi et al. 2014). Moderate inspiral times ($\sim \text{Gyr}$) produce the closest match between the phenomenological and GW-only models but still show a slight turnover in addition to an amplitude $\sim 10\%$ lower at all frequencies.

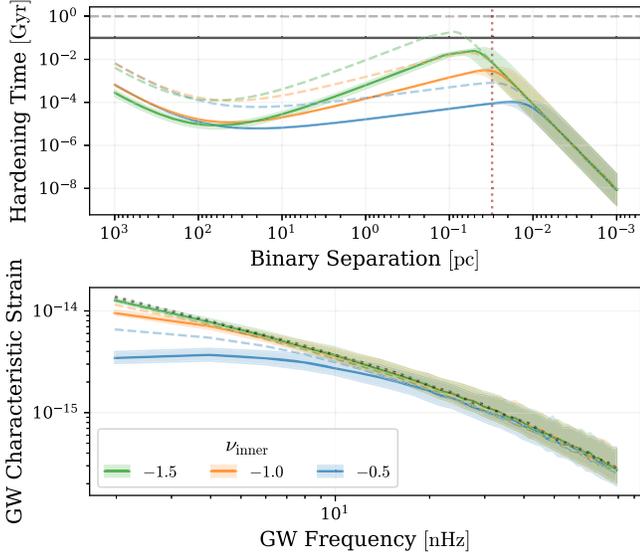


Figure 3. The SMBH binary hardening timescales and GW spectra are shown for varying values of the phenomenological binary evolution model parameters. The top panel shows the hardening timescale ($\tau = dt/d \ln a$), with the black horizontal lines corresponding to total binary lifetimes. Solid lines correspond to binary lifetimes $\tau_f = 0.1$ Gyr, while dashed are 1.0 Gyr. The small-separation hardening rate power-law index, ν_{inner} , is also varied, which changes the amount of time binaries spend at subparsec separations. The vertical dotted line shows the separation at which an $M = 10^9 M_\odot$ system reaches $f = 1/16.03$ yr. The bottom panel shows the resulting GWB characteristic strain. The dotted line shows GW-only driven evolution. For clarity, the top panel includes only binaries with $3 \times 10^8 M_\odot < M \leq 3 \times 10^9 M_\odot$ and $0.1 < q \leq 1.0$. In both panels, the shaded regions denote the interquartile range. The other parameters of these populations are the fiducial values for the *Phenom* set of models (Section 3.3). Variation in the binary evolution parameters significantly impacts the shape and low-frequency amplitude of the GWB spectrum.

3.4. Interpolation of Population Synthesis Models with GPs

In order to infer the properties of SMBH binary populations that are consistent with the GWB, we need to compare the theoretically expected GWB spectra from *holodeck* with the observed NANOGrav data. Previous work used analytic expressions for this, e.g., by fitting the GWB spectra with a single (for a population of circular binaries purely driven by GWs) or broken (to capture the turnover produced by environmental interactions; Sampson et al. 2015) power law. However, the properties of the SMBH binary population are only indirectly extracted from these fits, and disentangling potential covariances between population parameters is challenging. To overcome this limitation, Taylor et al. (2017) developed a modeling framework that directly links the properties of the GWB spectrum to the binary population parameters by training GPs on simulated GWB spectra from population synthesis models. Here we adopt this approach to interpolate the strain of the GWB across simulated *holodeck* libraries generated in discrete points of the binary parameter space to accurately predict the GWB spectrum at any point in the space.

The GPs provide a powerful interpolation method that parameterizes noisy data in terms of a multivariate Gaussian distribution with a mean vector and covariance function (see Aigrain & Foreman-Mackey 2022, for a review). The covariance functions can be custom built from a suite of versatile kernel functions allowing for quick adaptability to a variety of complex parameter spaces. While GPs are not sparse and lose efficiency in high-dimensional spaces (e.g., greater than a few dozen), one key

advantage of GP regression is that it provides an estimate of the uncertainty in the interpolation process (i.e., the prediction is probabilistic). Importantly, one can use this in an iterative process to adapt and improve the fitting. Additionally, the GP uncertainty can be propagated forward to our final statistics, allowing for a full marginalization over the interpolation uncertainties.

The GPs are trained on *holodeck* GWB spectra using the George GP regression library (Ambikasaran et al. 2015a) as in Taylor et al. (2017) and Arzoumanian et al. (2018). To capture fluctuations that arise directly from the discrete nature of the binary population, we train the GPs at each sampling frequency of the GW spectrum, f_i . GP regression assumes that the interpolated quantity (here the strain of the GWB) is smooth with respect to the interpolation variables (here the model parameters). The use of an independent GP at each frequency thus enforces smoothness in the GW spectrum across model parameter space at a given frequency but not across frequencies. Because the binary population is independent at each frequency, smoothness across frequencies is not expected in general. Two separate GPs are trained per frequency, one on the median value of $\log_{10}(h_c^2(f_i))$ and one on its standard deviation. This allows us to predict both the typical value and the typical spread of the strain and to account for the uncertainty in each value’s interpolation separately.

We select the training set (i.e., the library generation points that make up our model grid) from our multidimensional parameter space using Latin hypercube (LHC) sampling (e.g., see Taylor & Gerosa 2018, and references therein).⁸⁰ This offers an efficient method to generate a near-random set of parameter values, representative of the entire parameter space, with relatively few points. Since we aim to explore high-dimensional spaces, this type of sampling is necessary to keep the total number of simulations computationally tractable.

The training of the GPs proceeds as follows. Using the LHC method, we draw s samples in the binary parameter space. For each sample, we produce r realizations of the GWB spectrum using *holodeck* and calculate the median and standard deviation of $\log_{10}(h_c^2)$ at each frequency f_i . These means and standard deviations constitute the inputs for the training of the two GPs. For each point in the training set, the GPs require the value of the quantity on which they are trained (here the median or standard deviation of $\log_{10}(h_c^2(f_i))$) and, optionally, its uncertainty. Including uncertainties on the input values helps to avoid overfitting. We adopt the standard uncertainties for the sample mean and standard deviation. When training on the median, we estimate the uncertainty as the standard deviation divided by $r^{1/2}$; for training on the standard deviation, the uncertainty is given by the standard deviation divided by $[2(r-1)]^{1/2}$.

To test the performance of the GPs, we create a validation set with points in the SMBH binary parameter space that were not included in the training set. For each validation point, we calculate the median and standard deviation of $\log_{10}(h_c^2)$, both with GPs and with *holodeck* simulations. For comparison purposes, we label the value obtained from GPs as the “predicted” value, while the *holodeck* values are considered to be the “true” value. Based on this, an error (i.e., predicted minus true) can be calculated for the GP interpolation performance.

⁸⁰ One-dimensional LHC sampling divides the cumulative density function into a number of equal partitions and then chooses a random data point in each partition. Sample points in multiple dimensions are randomly combined. This approach ensures coverage of the domain, similar to a uniform grid, while not wasting samples at identically placed grid edges.

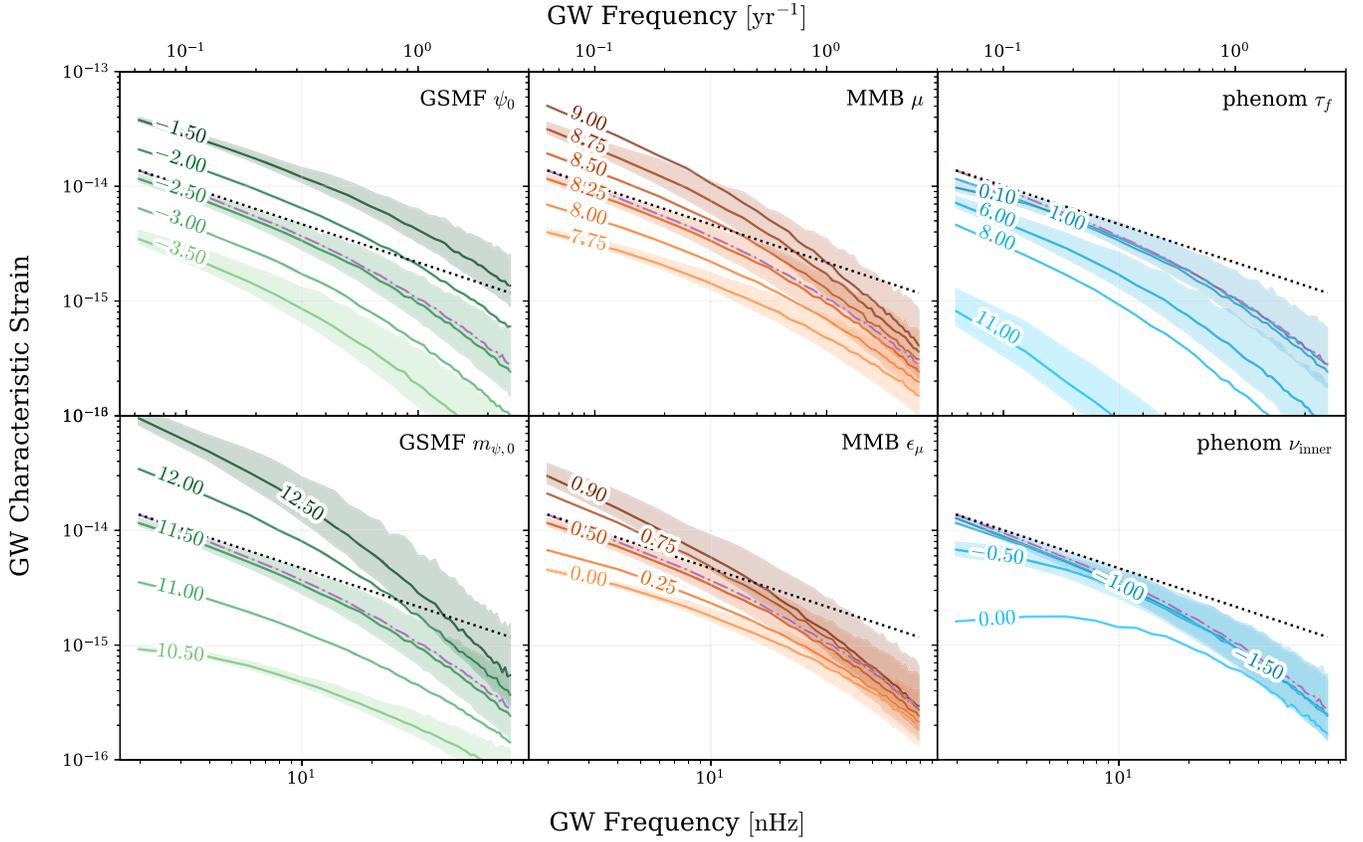


Figure 4. Comparison of GWB spectra for systematic variations in the parameters of our fiducial range of models (*Phenom*). In each panel, the indicated parameter is varied, while the other parameters are fixed at typical values. The purple dotted–dashed line shows the spectrum from a fiducial GW-only model, while the black dotted line shows a pure $f^{-2/3}$ power law as a reference. Variations in the GSMF parameters (ψ_0 , the dimensionless GSMF turnover mass) are shown in green. Variations in the $M_{\text{BH}}-M_{\text{bulge}}$ parameters (μ , the dimensionless $M_{\text{BH}}-M_{\text{bulge}}$ normalization, and ϵ_μ , the $M_{\text{BH}}-M_{\text{bulge}}$ scatter in dex) are shown in orange. Variations in the SMBH binary lifetime (τ_f in units of gigayears) and hardening power-law index (ν_{inner}) are shown in blue. The numeric label indicates the value of the parameter for that particular curve, and the fixed values for all other parameters are $\psi_0 = -2.5$, $m_{\psi,0} = 11.5$, $\mu = 8.25$, $\epsilon_\mu = 0.5$ dex, $\tau_f = 1.0$ Gyr, and $\nu_{\text{inner}} = -1.0$. The solid lines show the median for each parameter value from 10,000 realizations. The shaded regions indicate the 16th and 84th percentiles of the distribution. For clarity, we only plot the shaded regions for every other parameter value. Apart from the mostly degenerate mass parameters μ and $m_{\psi,0}$, which nonetheless have a significant influence on the GWB spectrum, each of these parameters impacts the GWB spectra in distinct ways. This indicates the promise of GW observations for constraining the SMBH binary population.

We used this approach to test a variety of kernels (i.e., covariance functions) along different directions in parameter space to determine which combination most accurately captured each parameter’s response to changing the GWB. We determined that two types of kernels were necessary: rational quadratic kernels for the phenomenological timescale τ_f and hardening power-law index ν_{inner} and squared exponential kernels for the remaining parameters. An iterative process of checking the performance of the GPs was used to determine the necessary number of LHC sample points and *holodeck* realizations to converge on a sufficient accuracy level. The performance of the GP trained on the median values is more sensitive to the choice of the number of sample points, while the performance of the GP trained on the standard deviations is more sensitive to the choice of the number of realizations. We found that training on $s = 2000$ LHC samples with $r = 2000$ *holodeck* realizations at each sample point was more than enough to acquire the desired accuracy level. Figure 6 shows the response of GPs trained on the *Phenom* library for the five frequencies used in our analysis. The reconstruction is quite accurate, with 99.4% (98.5%) of the test set cases for medians (standard deviations) falling within 10% of the actual value—significantly smaller than the standard deviation across spectral realizations.

3.5. Fitting Simulated GWB Spectra to PTA Observations

In Arzoumanian et al. (2018), once the GPs were trained, they were inserted into the full PTA likelihood calculation in order to obtain posteriors on the SMBH binary population parameters. As PTA data sets have grown in size, and with the new discovery of HD correlations, the likelihood computation time has increased. As such, inserting two GPs into the 15 yr data set’s likelihood calculation in order to obtain posteriors for *Phenom* was not an efficient analysis approach.

Instead, we use the *ceffyl* package (Lamb et al. 2023) to fit the interpolated GWB spectra to the previously computed free-spectrum posteriors of the cross-correlated timing-residual PSD. Fitting on intermediate PTA analysis products, such as the free-spectrum posteriors, offers a substantial speed-up by factors of 10^2-10^4 compared to directly fitting the full likelihood of the timing residuals. Importantly, the resulting posterior distributions of the GW spectral model parameters achieved by *ceffyl* have been found to be nearly identical to those obtained from the full likelihood approach.

In detail, we expand the likelihood, $\mathcal{L}(\mathbf{d}|\Theta)$, where \mathbf{d} is the PTA data (e.g., the TOAs) and Θ are the SMBH binary population parameters (e.g., the parameters from *Phenom*), by inserting an intermediate data product, such as the free-

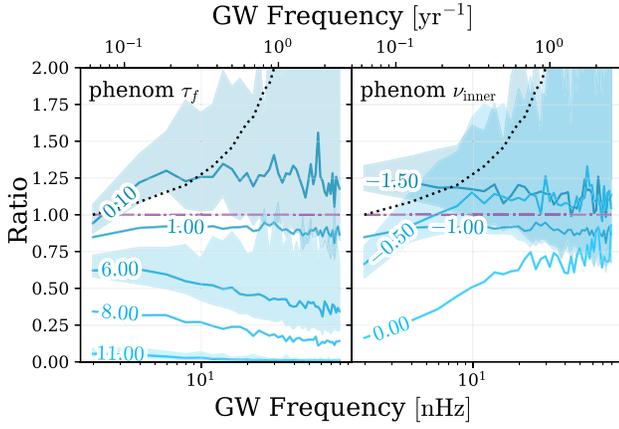


Figure 5. Comparison of GW spectra for systematic variations in two of the six model parameters (τ_f and ν_{inner}) normalized to a model in which binaries evolve only because of GW emission. Colors and line styles are as in Figure 4. The phenomenological hardening models show substantial deviations from the GW-only hardening models, especially for parameters far from our fiducial values of $\tau_f = 1$ Gyr and $\nu_{\text{inner}} = -1$. In particular, the spectra are suppressed at the lowest frequencies for larger values of ν_{inner} and produce a more pronounced turnover.

spectrum posteriors ($\log_{10}(\rho_i)$). Then, instead of directly calculating the fit of a GWB spectrum (generated by the trained GPs for a given draw of SMBH binary population parameters) to the TOAs, we compute the probability that a given GWB spectrum is supported by the free-spectrum posteriors. The expanded likelihood function is now given by

$$\mathcal{L}(\mathbf{d}|\Theta) \propto \prod_{i=1}^{N_f} \int d(\log_{10} \rho_i) p(\log_{10} \rho_i | \mathbf{d}) p(\log_{10} \rho_i | \Theta), \quad (28)$$

where N_f is the number of Fourier components used in the GWB analysis (5 or 14; but see also Section 2 and NG15gwb); $p(\log_{10} \rho_i | \mathbf{d})$ is the posterior probability density of $\log_{10}(\rho_i)$ (i.e., the free-spectrum posteriors), which are represented by highly optimized kernel density estimators; and $p(\log_{10} \rho_i | \Theta)$ is the probability of $\log_{10} \rho_i$ given a GWB spectrum from the trained GPs. Since the GPs are trained on the median and standard deviation of the characteristic strain $\log_{10}(h_c^2)$, they provide the mean and variance of a Gaussian when calculating $p(\log_{10} \rho_i | \Theta)$. The above likelihood is sampled through MCMC techniques to obtain the resultant posteriors on Θ .

While all of the libraries generated for GP training draw uniformly from the SMBH binary population parameter space, when we perform the MCMC analysis, we have the opportunity to place different priors on each parameter. For the analysis in this paper, we utilize two distinct prior setups: a uniform prior and a set of astrophysical priors based on galaxy observations (e.g., see Table B1). When relevant, we denote the prior distribution shape in combination with the library designation as, e.g., *Phenom+Uniform* or *Phenom+Astro* (see Table B2).

4. Results

We simulate populations of SMBH binaries using phenomenological (*Phenom*) and GW-only (*GWOnly*) models. We create `holodeck` libraries of GWB spectra at fixed points of the SMBH binary parameter space and interpolate them with GPs. We fit the models to the 15 yr free-spectrum posteriors

considering the *HD-w/MP+DP+CURN* as the fiducial 15 yr NANOGrav results for this analysis (but we also fit the *HD-DMGP* posteriors for comparison) using both uniform and astrophysically motivated priors (see Table B1). As shown in Table B2, the *Phenom* library is fit against the data using both uniform and astrophysically informed priors (*Phenom+Uniform* and *Phenom+Astro*), while the *GWOnly* library is fit only with uniform priors (*GWOnly+Uniform*). Our results are summarized as follows.

In all of our analyses, we find that the NANOGrav 15 yr data set is consistent with a GWB produced by a population of SMBH binaries. In the first, most simplified approach, power-law fits⁸¹ to both the observed GWB spectrum and those from simulations produce amplitudes and spectral indices that overlap in the 2σ and 3σ regions, depending on the model (Section 4.1). The remainder of this section presents the results of our systematic approach of fitting simulated SMBH binary populations to the data, which yield more realistic GWB spectra that match the 15 yr results (Section 4.2). From these fits, we obtain posterior distributions on the uncertain astrophysical parameters of the SMBH population synthesis models (Section 4.3) and make predictions from our models for the properties of the population of SMBH binaries that produce the GW observations (Section 4.4).

4.1. Comparison of Idealized Power-law Fits to GWB Spectra

The approach of fitting simple power-law models to the GWB is a common one in the literature. While idealized power-law fits to GWB spectra neglect most of the information imprinted by astrophysical processes on the background, they are effective in broadly examining the consistency between simulated binary populations and PTA data sets. Therefore, we carry out this straightforward analysis as a first check of the *Phenom* and *GWOnly* libraries before implementing the full methodology described in Sections 3.4 and 3.5. In practice, we constrain the amplitude, A , and slope, γ , of an idealized power-law GWB spectrum (in the timing-residual PSD; Equation (12)) with a nonlinear least-squares fit to the GWB spectra from each realization of the binary population from the *Phenom* and *GWOnly* libraries using the five lowest frequency bins. We then compare these to the results of power-law fits of the 15 yr data, illustrating their overlap in the $A - \gamma$ parameter space.

In Figure 7, we show the range of GWB amplitudes and spectral indices (in the timing-residual PSD) based on these fits. We see that the amplitudes and power-law indices vary significantly across the simulated GWB spectra. Even for the *GWOnly* models, which match the premise of the analytic $\gamma = 13/3$ ($\alpha = 2/3$) models, our simulations yield indices typically varying from 4 to 5.5 in the 95% credible region. Recall that even GW-only binary evolution with circular orbits does not produce a pure power-law spectrum owing to the steepening of the spectrum at higher frequencies, where the finite number of binaries in each frequency bin becomes important. The slight offset of the *GWOnly* models toward steeper values of $\gamma > 13/3$ reflects this higher-frequency spectral steepening caused by finite-number effects. The *Phenom* libraries, which self-consistently model the effects of

⁸¹ Note that NANOGrav constraints are derived primarily at lower frequencies. Fitting power laws and extrapolating the amplitudes to $f = (1 \text{ yr})^{-1}$ can lead to amplitudes that differ more significantly at this frequency than at $f = (10 \text{ yr})^{-1}$, for example. See Appendix A.

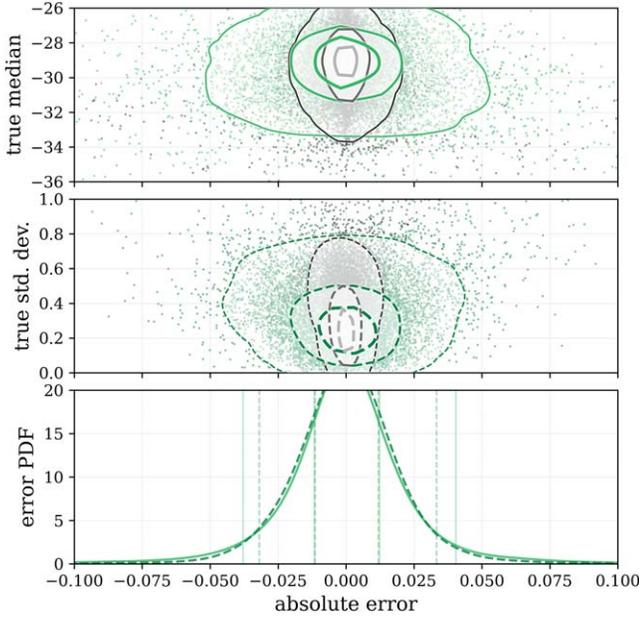


Figure 6. Accuracy of trained GP interpolants for our fiducial six-dimensional parameter space *Phenom*, aggregated over the first five frequency bins used in our primary analyses. The top panel shows the error (\equiv predicted – true) on the median $\log_{10} h_c^2$ (solid lines) for the training set (gray) and validation set (green), while the middle panel shows the error on the standard deviation (dashed lines). Contours contain 20%, 50%, and 90% of the population. The bottom panel shows the distributions of errors for both validation sets. The vertical lines bound the 50% and 90% region of the errors. The GPs predict 99.4% (98.5%) of the validation set for medians (standard deviations) within 10% of the actual value. Both the training and validation sets contain 2000 sample points, with medians and standard deviations calculated over 2000 realizations, which we have found to be more than sufficient in accurately training the GPs.

environmental interactions on binary evolution, produce much wider ranges of spectral indices, as disparate as 2–7, with lower values corresponding to shallower characteristic strain spectra (i.e., increasing across the five lowest frequency bins). Note that we exclude from Figure 7 a small number ($\sim 1\%$) of *Phenom* samples in which all binaries stall and thus produce zero GWB.

We also show the power-law parameter posteriors for the fiducial *HD-w/MP+DP+CURN* free-spectrum posteriors and the *HD-DMGP*, which we use for comparison. While the 15 yr free-spectrum posteriors are not perfectly fit by power-law models (see Figure 1), the differences between these two models highlight that the measurement of a spectral index is particularly sensitive to choices of fit in the 15 yr data and features in particular frequency bins (NG15gwb).

We use the $A - \gamma$ fits as a general measure of parameter space coverage. Figure 7 demonstrates that the range of simulated populations is able to reproduce the measured GWB within the 2σ curve of the *Phenom* library and between the 2σ and 3σ curves of the *GWOnly* library.

4.2. The GWB Is Consistent with Expectations from Populations of SMBH Binaries

The consistency between the 15 yr NANOGrav data set and GWBs produced by SMBH binaries is best supported by an analysis of the full range of astrophysical information contained in the free-spectrum posteriors. Fitting the GPs trained on the *GWOnly* and *Phenom* libraries to the 15 yr data

(with uniform and astrophysically motivated priors) facilitates a comparison of observations to the GWB spectra from SMBH binary populations that is agnostic to any particular spectral model (including a power law).

Figure 8 shows GWB spectra produced by our simulated SMBH binary populations that accurately fit the 15 yr *HD-w/MP+DP+CURN* free spectrum. As mentioned above, the *Phenom* library is fit with both uniform and astrophysically informed priors (*Phenom+Uniform* and *Phenom+Astro*), while the *GWOnly* library is fit only with uniform priors (*GWOnly+Uniform*). Thin curves show 200 random draws of the binary parameter posterior distributions for each of the above models, with thick lines denoting the maximum-likelihood spectra for each model.

Both libraries are able to fit the GWB within the 15 yr posteriors. However, the *GWOnly* spectra have more difficulty matching the data, as indicated by their preference for the edges of the 15 yr free-spectrum posteriors in the highest and lowest frequency bins and the best-fit spectrum missing the highest-probability regions of the 15 yr GWB data. As a comparison, power-law fits are shown for the idealized $\gamma = 13/3$ ($\alpha = 2/3$) spectral indices obtained from analytic calculations of SMBH binaries (Phinney 2001). The *GWOnly* models, which more closely resemble these analytic estimates, tend to be steeper than the bulk of the 15 yr distributions. In contrast, the maximum-likelihood spectra and likelihood draws from the *Phenom* model exhibit noticeable spectral turnovers to match the 15 yr data. While these results are suggestive of a low-frequency turnover or flattened spectrum, they are still consistent with an $\alpha = 2/3$ power law and the associated GW-driven evolution.

4.3. Parametric Constraints on SMBH Binary Models

The MCMC exploration of the likelihood in Equation (28) returns constraints (posterior distributions) on the parameters of the population synthesis models based on the observed GWB spectrum. The peaks of the marginalized posteriors indicate the most likely values of the parameter space that the binary population must occupy in order to produce the 15 yr free-spectrum posteriors. Figure 9 shows the posteriors of these binary population parameters for the *Phenom* binary evolution model. Results are compared for different prior choices, i.e., the *Phenom+Uniform* and *Phenom+Astro* fits. Owing to the substantial uncertainty in the GWB spectrum at NANOGrav’s current sensitivity, the posteriors are sensitive to the assumed priors, and only weak parameter constraints can be made.

However, we can identify some general trends among the preferred parameter values. The measured amplitude of the GWB strongly prefers a combination of efficient mergers occurring in high-mass systems. The data favor short binary lifetimes (τ_j), high GSMF number densities (ψ_0), and high characteristic masses ($m_{\psi,0}$, μ). It is worth noting that the range of priors in the *Phenom+Uniform* fits is quite wide compared to the typical values adopted in the astronomical literature (see references with Table B1). While the parameters in the *Phenom+Astro* model are more constrained, the posteriors are still fairly broad. Because our models utilize simplified analytic prescriptions for each physical component, we use broader parameter distributions in *Phenom+Uniform*, in part to introduce some added flexibility. Nonetheless, Figure 9 demonstrates that the posteriors almost uniformly favor parameters that produce larger GWB amplitudes (e.g., see also Figure 4). This suggests that the

amplitude of the GWB inferred from the 15 yr data set is difficult to reach with the standard values of some astrophysical parameters.

Very long binary lifetimes are disfavored for both *Phenom+Uniform* and *Phenom+Astro*. A large fraction of binaries with such long lifetimes would fail to reach the NANOGrav frequency band, resulting in lower GWB amplitudes, inconsistent with the GW data. Flatter values of the hardening rate power-law index, $\nu_{\text{inner}} \gtrsim -1.0$, are also preferred, as they produce spectral turnovers in the lower frequency bins (see Figure 3) resembling what is seen in the 15 yr data. Steeper values of ν_{inner} correspond to binaries that spend more time at $\sim 10^{-2}$ –1 pc separations and transition into the GW-dominated regime earlier, at frequencies below the PTA band. Flatter values of ν_{inner} correspond to very efficient inspiral through this range of separations, leading to environmentally driven evolution even in the lower PTA band that produces noticeable GWB attenuation and a sharp low-frequency spectral turnover (see Figure 5). We note that this is dependent on the parameterization of our binary evolution model and the parameters varied in the *Phenom* library. Steeper evolution profiles at very large separations (\sim kpc), i.e., larger values of ν_{outer} , could similarly produce low-frequency turnovers (but this is not explored in this paper, since ν_{outer} is kept fixed throughout). In either case, efficient binary inspiral in the environmentally driven regime produces the most noticeable spectral turnovers.

The posteriors for the GSMF and $M_{\text{BH}}-M_{\text{bulge}}$ parameters differ noticeably for uniform versus astrophysical priors, unlike the binary inspiral parameters. The posteriors for the normalization and characteristic mass of the GSMF (ψ_0 and $m_{\psi,0}$) favor values at the higher end of the prior range, especially in the *Phenom+Uniform* fits, in which the galaxy number densities are pushed against the edges of the prior. We also see higher values of these posteriors when binary lifetimes are longer, such that larger fractions of binaries stall before reaching the PTA band.

Unsurprisingly, the GSMF characteristic mass ($m_{\psi,0}$) is almost entirely degenerate with the $M_{\text{BH}}-M_{\text{bulge}}$ mass normalization (μ), as indicated by the diagonal band in the respective two-dimensional posterior. The scatter in the $M_{\text{BH}}-M_{\text{bulge}}$ relationship (ϵ_μ), however, shows different trends. This is likely due to two factors. First, increasing ϵ_μ primarily increases the GWB amplitude in the lower frequency bins (Figure 4), as larger scatter preferentially increases SMBH masses, and higher masses are more prevalent at lower frequencies (discussed more below). Second, larger values of ϵ_μ also produce significant variance across multiple population realizations, which may decrease the aggregated likelihoods when calculating fits.

In Figure 10, we compare the one-dimensional distributions of parameter priors versus posteriors for the *Phenom+Uniform*, *Phenom+Astro*, and *GWOnly+Uniform* models. Particularly in the case of the $M_{\text{BH}}-M_{\text{bulge}}$ parameters, we see that the posteriors closely follow the priors. While still consistent with the priors, the GSMF parameters are pushed noticeably towards higher values, even for the *Phenom+Astro* fits.

Figure 10 also shows fits using the *GWOnly* library. Note that this library does not include the τ_f or ν_{inner} parameters by definition. The posterior distributions for the GSMF and $M_{\text{BH}}-M_{\text{bulge}}$ parameters are generally consistent in both the *Phenom* and *GWOnly* models, with only weak constraints.

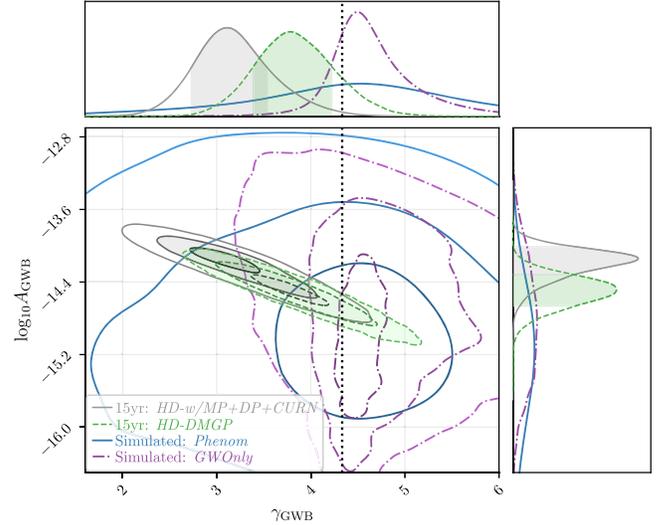


Figure 7. Power-law amplitude (A) and spectral index (γ) from purely power-law fits to HD free-spectrum model posteriors from the 15 yr data set compared to simulated GWB spectra from the *holodeck* libraries. Data set fits include both the 15 yr *HD-w/MP+DP+CURN* and *HD-DMGP* models for comparison. Fits to the five lowest frequency bins of spectra from two *holodeck* libraries are shown: the self-consistent phenomenological binary evolution model (*Phenom*) and the purely GW-driven evolution model (*GWOnly*). We show 1σ , 2σ , and 3σ contours for each. The analytic GWB PSD power-law index of $\gamma = 13/3$ is shown as a reference (black dashed line). The spectral shape of the HD signal present in the 15 yr data set is broadly consistent with expectations for a GWB from a binary SMBH population. The amplitude is toward the higher end of the predictions, and the recovered spectral index deviates from the idealized power law in similar ways as the phenomenological binary evolution model.

While *GWOnly* shows the same preference for high values of ψ_0 as the phenomenological model, the preference is less pronounced. This is likely due to the decrease in GWB power at the lowest frequencies when a spectral turnover is induced, which is consistent with the covariance seen in Figure 9 between binary lifetime and GSMF normalization.

It is important to note that these parametric constraints must be interpreted in the context of the semianalytic binary evolution models used to generate the binary populations and corresponding GWB spectra. For example, the usage of a fixed-time phenomenological binary evolution model is forcing a particular relationship between typical binary masses and the degree of low-frequency spectral turnover. Another model, in which the degree of environmental coupling scales differently with binary mass (or, similarly, host galaxy properties; e.g., Kelley et al. 2017a), may produce different dependencies and thus different posteriors. We are also assuming a fixed $M_{\text{BH}}-M_{\text{bulge}}$ relationship for all redshifts, while the canonical $M_{\text{BH}}-M_{\text{bulge}}$ relationship in the literature is specifically calibrated to the local Universe. Our values of μ and ϵ_μ for both the *Phenom* and *GWOnly* models should thus be interpreted as “redshift-averaged” quantities.

4.4. Inferred Properties of the SMBH Binary Populations

While a large amount of information is encoded in the GWB spectra, there are numerous degeneracies, particularly in the current low signal-to-noise regime. For example, given a particular GWB spectral shape, a certain GWB amplitude can be produced by a large number of lower-mass SMBH binaries or a small number of higher-mass SMBH binaries. To determine the characteristic properties of SMBH binaries

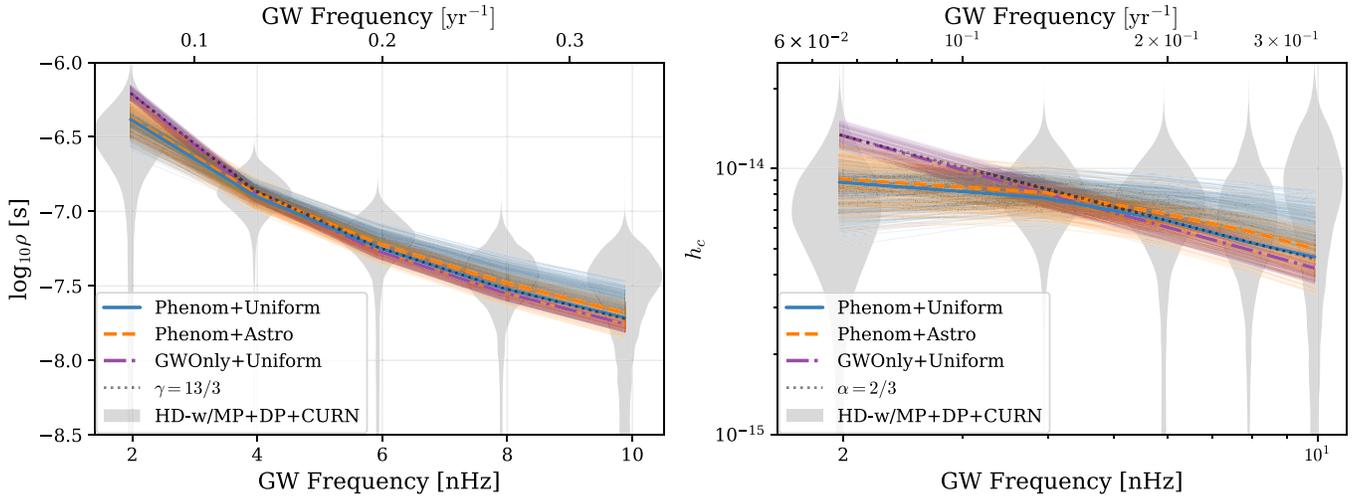


Figure 8. The GWB spectra from simulated SMBH binary populations that best fit the 15 yr free-spectrum data. The left panel shows the square root of the cross-correlated timing-residual power (ρ), and the right panel shows the characteristic strain (h_c). The GP-interpolated spectra are shown, with thick lines showing maximum likelihoods and thin lines showing 200 random draws from the posteriors. The *Phenom* library uses self-consistent binary evolution models, while *GWOnly* assumes purely GW-driven evolution. For the former, fits using uniform priors (*Phenom+Uniform*) are compared against more informed, astrophysically motivated ones (*Phenom+Astro*). Power-law fits to the 15 yr spectra with $A_{\text{yr}} = 2.1 \times 10^{-15}$ are also shown as dotted black lines.

contributing to the GWB, we calculate the distribution of GW-weighted average binary parameters. We use 1000 draws from the posteriors of the *Phenom+Uniform* fits. For each draw, the h_c^2 -weighted parameters are calculated over 100 realizations. This gives a distribution of the average parameters for each draw and realization, which are plotted in Figure 11. As in all of our analyses, we fit binary population models to only the five lowest frequency bins in the 15 yr data set. However, in order to better visualize the trends in binary population parameters with GW frequency, Figure 11 shows the *Phenom+Uniform* library priors and posteriors for 10 frequency bins.

The GWB is characterized by the most massive SMBH binaries in the Universe, with $M \gtrsim 10^{8.5} M_\odot$, and extending to just above $10^{10} M_\odot$ at the lowest frequency bins. At higher frequencies, as binaries evolve more quickly and fewer binaries occupy each frequency bin, these most massive systems become rarer, and the typical masses decrease. Because of the trend in mass, the typical separations of binaries decrease more rapidly than $f^{-2/3}$, as would be expected for a fixed mass. The binary total masses are the most strongly constrained parameters when comparing between the library priors and the posteriors. This is unsurprising given (a) the strong dependence of GW strain on binary mass and (b) the numerous varying model parameters that affect the masses. The mass-ratio distributions, on the other hand, are nearly constant across the band and narrowly localized for both the priors and posteriors. Typical binary mass ratios are almost entirely above $q \sim 0.5$. Note that this is determined primarily by our fiducial parameters for mass-ratio dependence in the GPF and GMT.⁸² For the latter in particular, the GMT scales as q^{-1} , which strongly disfavors extreme mass-ratio pairs.

Across the PTA band, the binaries producing the GWB are typically at many hundreds to a few thousand comoving megaparsecs ($z \approx 0.15\text{--}0.9$). The average redshift posteriors are higher than the priors due to fits preferring shorter binary

lifetimes. Binary separations are tightly constrained by the strong constraints on the binary masses. Typical separations are just below 10^{-1} pc at the lowest frequency bin (≈ 2 nHz), down to just below 10^{-2} pc at the 10th frequency bin (≈ 20 nHz). Projecting these separations at the cosmological distances of the binaries leads to angular separations of tens of microarcseconds.

Having explored the distributions of GWB-weighted binary properties from our fiducial model (Figure 11), we now examine which binary parameter ranges contribute most to the GWB signal. In Figure 12, we show the fraction of the GWB contributed by different portions of the binary population. Using the same posteriors as were sampled previously, we plot the fraction of h_c^2 that is contributed by binaries at each binary parameter value (solid lines). The results are also separated by GW frequency (colors). We compare these GWB fractions to the total number of binaries in our simulated populations that are emitting at a given frequency (dashed lines). There is a stark difference between the number of binaries emitting at each frequency as a function of M_{tot} and their relative contribution to the GWB signal. Lower-mass binaries are far more common, but the signal is dominated by the rare massive black holes near $M_{\text{tot}} = 10^{9.5}\text{--}10^{10} M_\odot$. For example, binaries between $M_{\text{tot}} = 10^{9.2}$ and $10^{10.4} M_\odot$ make up 2.6×10^{-4} of all binaries emitting in the lowest frequency band, but they make up 73% of the signal at that frequency. Similarly, we are preferentially sensitive to the largest mass ratios.

Although the bulk of the GWB signal is made up of binaries at $z > 0.4$, it is most sensitive to the nearest binaries relative to the underlying SMBH binary population. Because, all else being equal, the binaries that come from later galaxy mergers (lower z_{init}) will enter the PTA band at later times (lower z_{final}), the GWB source population is biased toward these lower-redshift sources. These biases affect all frequencies relatively uniformly, with the mass bias being slightly more pronounced at lower frequencies. This reflects the increased total number of binaries at low frequencies, which allows for the rarer, higher-mass SMBH binaries to dominate, as also seen in Figure 11.

Note the distinction between Figure 11, which shows h_c^2 -weighted binary parameter distributions, and Figure 12, which

⁸² We do allow the GPF mass-ratio dependence ($\gamma_{p,0}$) to vary in the *GWOnly-Ext* and *Phenom-Ext* libraries (see Appendix C). The parameter posteriors are virtually identical to the priors, suggesting that varying the mass-ratio dependence has little effect on the goodness of fit.

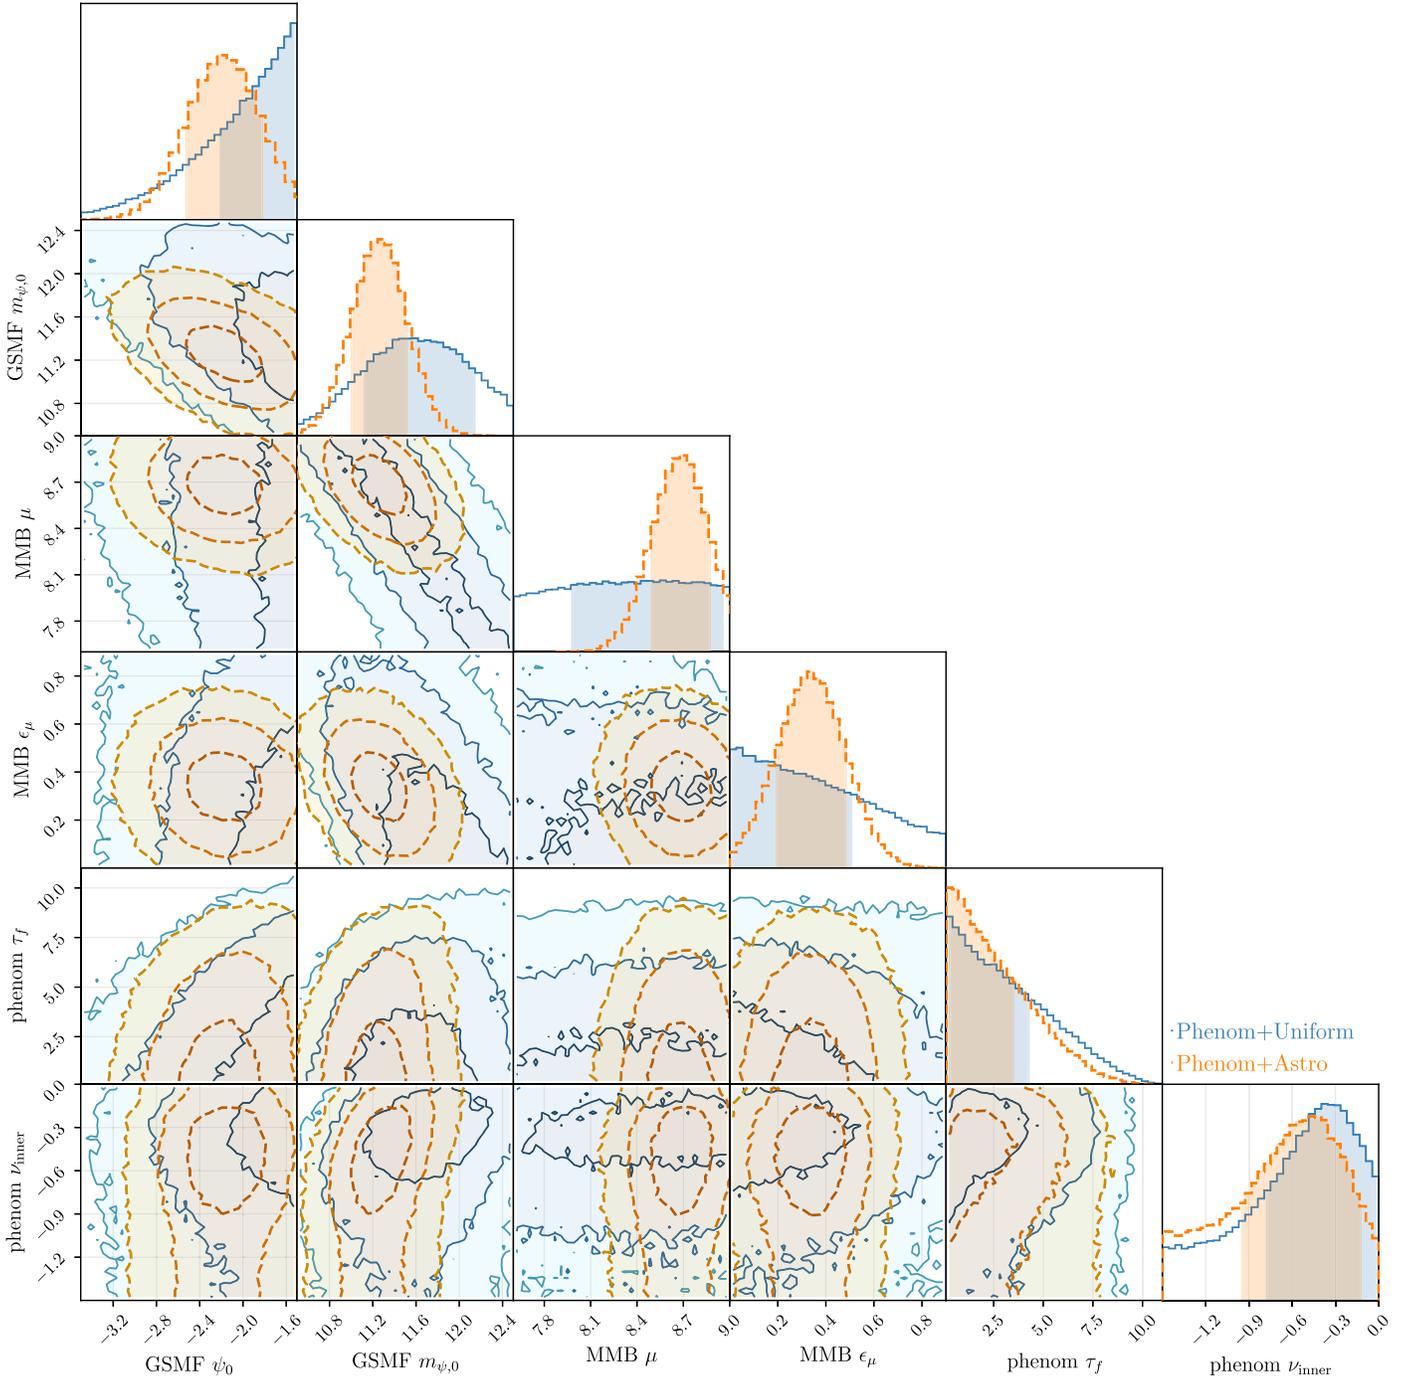


Figure 9. Binary evolution parameter posteriors from fitting against the 15 yr *HD-w/MP+DP+CURN* free spectrum. The parameters correspond to the phenomenological library *Phenom*: binary lifetime (τ_f), hardening power-law index (ν_{inner}), GSMF normalization (ψ_0) and characteristic mass ($m_{\psi,0}$), and $M_{\text{BH}}-M_{\text{bulge}}$ normalization (μ) and scatter (ϵ_μ). The fiducial *holodeck* spectral library is fit to the five lowest frequencies utilizing two separate priors. *Phenom+Uniform* uses flat priors across all parameter spaces (blue solid contours), while *Phenom+Astro* (orange dashed contours) uses constrained priors on the GSMF and $M_{\text{BH}}-M_{\text{bulge}}$ parameters that encapsulate observational measurements for these parameters. The contours in two dimensions correspond to the 1σ , 2σ , and 3σ regions, and the shaded regions in the one-dimensional plots are 1σ regions. Although individual parameters are only weakly constrained, the data strongly prefer efficient mergers in high-mass systems.

shows the fractional contribution to the GWB of binaries with given parameters. These figures are very closely related, but Figure 11 shows the representative properties of GWB binaries, while Figure 12 shows the fractional GWB contribution of actual binaries.

Our simulated populations also contain individually loud, high-mass binaries that can contribute substantially to the GWB. These sources, apparent in Figure 12 as spikes in the GWB fraction at high total masses, are likely the types of sources that will be detectable as continuous-wave signals by

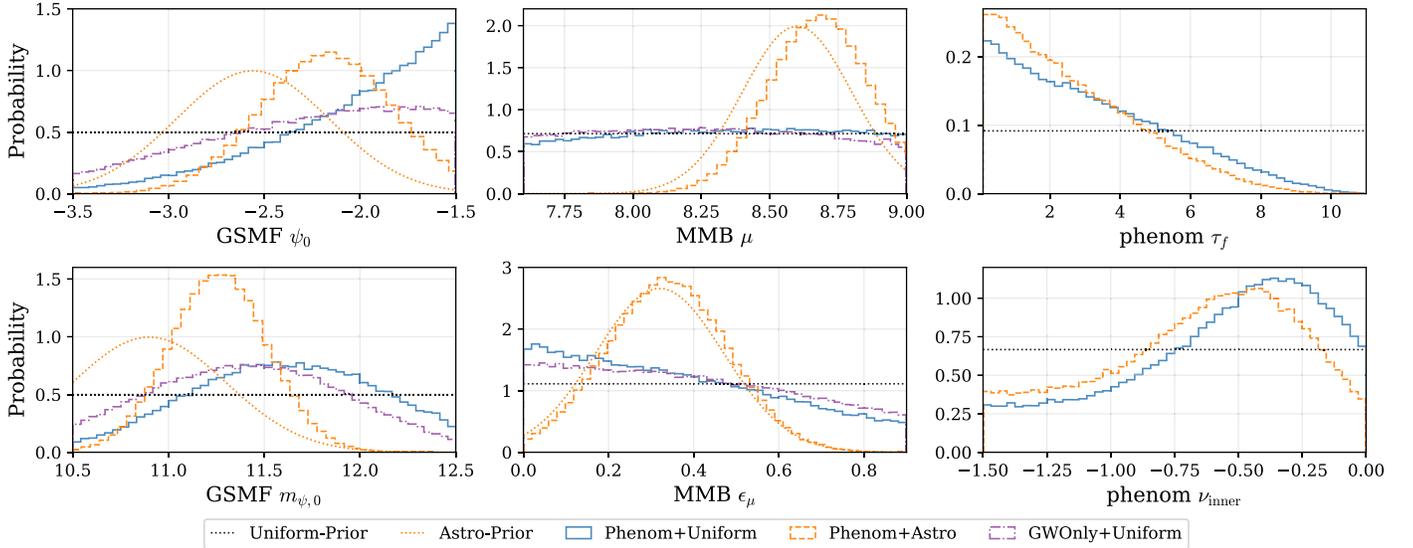


Figure 10. One-dimensional binary evolution parameter posteriors from fitting against the 15 yr *HD-w/MP+DP+CURN* free spectrum. Priors (dotted) and posteriors (solid) are shown for the *Phenom* model with both uniform and astrophysical priors. Note that the astrophysical and uniform priors are identical for τ_f and ν_{inner} (see Table B1). The $M_{\text{BH}}-M_{\text{bulge}}$ parameters are the least constrained (i.e., the posteriors largely resemble the priors), while the GSMF parameters are the most constrained. This is consistent with our expectations, given each parameter’s impact on the simulated GWB (see Figure 4). However, these one-dimensional projections can miss significant two-dimensional constraints due to various projection effects (e.g., the combined $M_{\text{BH}}-M_{\text{bulge}}$ mass normalization and $m_{\psi,0}$ two-dimensional space is much more constrained than each parameter individually, as seen in Figure 9). Overall, there is broad consistency between the posteriors for the *Phenom* models regardless of which prior is used, showing that these constraints are coming directly from the GWB spectrum itself and are not based solely on our binary population model construction. Distributions are also shown for the *GWOnly* model, which gives posteriors broadly consistent with the phenomenological models. However, they are peaked at slightly lower parameter values, indicating an overall lower amplitude for their simulated GWB. This is most likely due to this model’s rigidity, which does not allow for significant deviations from a power law-like GWB and leads to an inability to capture the behavior of the lowest frequency bin in the 15 yr free spectrum (see Figure 8).

PTAs (Agazie et al. 2023d; Arzoumanian et al. 2023). While they can occur at a range of frequencies, they are typically expected mid-band ($f \sim 3\text{--}30$ nHz), where the overall GWB amplitude has dropped somewhat but a sizable population of binaries remains (Kelley et al. 2018; Bécsy et al. 2022).

The dashed lines in the first panel of Figure 12 indicate the general shape of the mass function of binaries contributing to the GWB. In Figure 13, we explicitly calculate the implied mass function of SMBH binaries at frequencies within the PTA band for several redshifts (z_{final}). Distributions are shown for populations drawn from a sample of the *Phenom+Uniform* parameter space priors and from our 15 yr spectral fit posteriors. For comparison, we also show the mass functions for nonbinary SMBHs at each redshift, and in the last panel, we show the fraction of SMBHs at each redshift that are in binaries. At $z \geq 1.0$, the implied mass function is consistent with a wide range of values, indicating weak PTA constraints on high-redshift SMBH binary populations. This partly reflects the steep drop in binary fraction with increasing redshift; binaries that emit in PTA bands are $\sim 10\text{--}100$ times rarer at $z = 1.5$ than at $z = 0.25$.

Because of this, the binary number density increases over time (note that the nonbinary number density also increases with time as required for monotonic SMBH growth, but the evolution of the nonbinary mass function is slight compared to the binary mass function evolution). The shape of the binary mass function also significantly evolves, such that at lower redshifts, we see a much clearer turnover at $M \sim 10^9 M_{\odot}$ that more closely traces the shape of the nonbinary mass function. The binary mass function is also much more tightly constrained at low redshifts than at high redshifts, especially at the high-mass end. These $\gtrsim 10^9 M_{\odot}$ binaries at $z < 1$ are precisely the objects that comprise the bulk of the GWB (Figure 12).

Figures 11–13 examine the properties of binaries that emit in PTA bands. In Figure 14, we show the fraction of all binaries that reach the lowest frequency bin of the 15 yr data before redshift zero. Although binaries in the PTA bands will not reach coalescence on human timescales, the binary lifetime from PTA frequencies to merger is significantly shorter than a Hubble time.⁸³ We therefore use the fraction of binaries reaching PTA frequencies as a proxy for the fraction of systems that coalesce entirely before redshift zero.

Whether a given binary coalesces or stalls is determined by its formation redshift combined with the binary evolution time to reach the PTA band. Fitting to the 15 yr data strongly favors short binary lifetimes, which drives the difference between priors and posteriors. The redshift (bottom panel) at which the coalescing fraction reaches zero marks the redshift at which the look-back time of the Universe matches the binary lifetime of the model. The median posterior value of this lifetime is ≈ 2.8 Gyr, corresponding to a redshift of $z \approx 0.25$, where the median coalescing fraction reaches zero. The gradual increase of the coalescing fraction with redshift after this point is due to the additional delays from the GMTs. The coalescing fractions reach unity once the combined binary lifetimes and GMTs are longer than the look-back time.

The gradual increase in coalescing fraction with mass ratio is due primarily to the GMT’s strong dependence on the mass ratio ($T_{\text{gal-gal}} \sim q_*^{-1}$) in the fiducial model. The more gentle decline in the coalescing fraction at the highest total masses is due to more extreme mass-ratio systems, as seen in the comparison between the dashed ($q > 0.2$) and solid (all q) lines.

⁸³ For the binary masses characteristic of the GWB (e.g., Figure 11), the inspiral time from the lowest frequency bin ranges from roughly 0.05 to 50 Myr.

This is caused by a combination of increased GMTs and the increased binary coalescing times within the PTA band produced by more extreme mass-ratio systems that have longer GW-driven inspiral times.

We compare the *Phenom* model to a higher-dimensional library (*Phenom-Ext*) that includes variations in two of the GMT parameters, as well as others, in Appendix C and Figure C1. Generally, all of the recovered posteriors are consistent between the different libraries for the parameters that they have in common, suggesting that our choices of fiducial parameters are sufficiently representative of the binary evolution parameter space. The posteriors on the additional parameters themselves are generally broad. The exception is the GMT parameters, which, like the phenomenological evolution parameters, strongly favor shorter lifetimes as a way of producing higher GWB amplitudes.

5. Discussion

The NANOGrav PTA has detected a common-spectrum correlated stochastic process that is consistent with an astrophysical GWB. In our 15 yr data set (NG15, NG15detchar), we find evidence of the HD correlations that would definitively mark this signal as GW in origin (NG15gwb). In this paper, we have presented analyses of the NANOGrav 15 yr data set under the assumption that these data represent a GWB produced by SMBH binaries. With reasonable choices of astrophysical parameters governing galaxy masses, galaxy mergers, SMBH masses, and SMBH binary inspiral timescales, we are able to reproduce the inferred GWB amplitude and spectral shape. We find that the data are suggestive of a GWB spectral turnover at low frequencies, as expected for binary inspiral driven by astrophysical environments. However, the broad free-spectrum posteriors from the 15 yr data are still consistent with the canonical $\alpha = 2/3$ ($\gamma = 13/3$) power law expected for GW-driven inspiral.

Figure A1 compares the posteriors for the GWB amplitude and spectral index inferred from the 15 yr data with a wide variety of GWB model predictions in the literature (see also Table A1). Although the inferred GWB amplitudes are within the range of some of these model predictions, they lie at the high-amplitude end of this range. The implied GWB amplitude from the NANOGrav data therefore indicates that SMBH binary model parameters differ from standard expected values, although they still remain within reasonable bounds.

In this analysis, we have generated simulated populations of SMBH binaries and GWB spectra and fit them to the observed 15 yr signal. Our fiducial models explore a six-dimensional space of binary evolution parameters. Relative to the typically assumed values for these parameters, our results indicate that the inferred GWB amplitude could be achieved with short binary hardening timescales, higher galaxy number densities (translating to higher galaxy merger rates), or higher normalization of the $M_{\text{BH}}-M_{\text{bulge}}$ relation. This may be accomplished if multiple parameters differ somewhat from standard expectations or a small number of parameters differ more significantly.

Our models also demonstrate that the GWB signal is strongly dominated by the most massive, high mass ratio SMBH binaries, even among the subset of SMBH binaries emitting in PTA bands. The binaries contributing to the GWB form at typical redshifts of $z \approx 0.15-0.9$. Their typical separations (which are tightly constrained via the SMBH masses) range from ~ 0.1 to 0.01 pc; this corresponds to binary angular separations of tens of microarcseconds. Owing to the

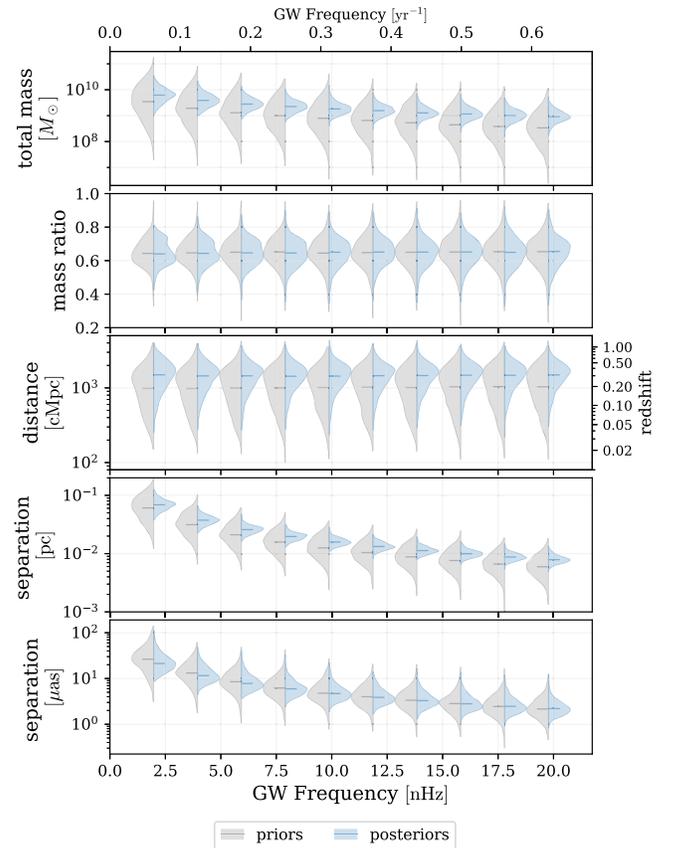


Figure 11. Distributions of GWB-weighted binary parameters versus GW frequency for populations drawn from the *Phenom+Uniform* library priors (left-hand violins in gray), as well as from the posteriors after fitting to the first five frequency bins of the 15 yr NANOGrav data (right-hand violins in blue). For each population sample, the h_c^2 -weighted averages of each parameter are calculated, and the distributions of those averages are plotted here. The GWB favors high total binary masses, especially in the lowest frequency bins, and typical binary separations range from $\sim 10^{-1}$ to 10^{-2} pc between 2 and 20 nHz.

short binary lifetimes preferred by the 15 yr data set, most of these binaries will merge by $z=0$; the coalescing fraction is near unity for binaries that form by $z \sim 0.25$. In addition, we note that our simulated binary populations contain loud, high-mass continuous-wave sources that could be detected above the GWB.

Because we are currently in the low signal-to-noise regime of GWB observations, we are still limited in our ability to make stringent parametric constraints. In this analysis, our constraints on the binary population inferred from the GWB spectral shape and amplitude are dependent on both our choice of priors and which 15 yr GWB measurements are used. NANOGrav continues to collect data from an ever-increasing number of pulsars. Forty-three pulsars were included in the 12.5 yr⁸⁴ analysis (Arzoumanian et al. 2020) versus 67 pulsars in the current 15 yr analysis. As of summer 2023, we were timing roughly 75 pulsars with a total baseline of over 17 yr. Also, NANOGrav data are currently being combined with those from other PTAs to create a new IPTA data set that will contain over 100 pulsars (Antoniadis et al. 2022). These efforts will improve our GWB measurement accuracy, along with our ability to constrain SMBH binary physics. The theoretical forecasts from Pol et al. (2021), for example, suggest that analyses using the

⁸⁴ See Alam et al. (2021) for the complete 12.5 yr data set.

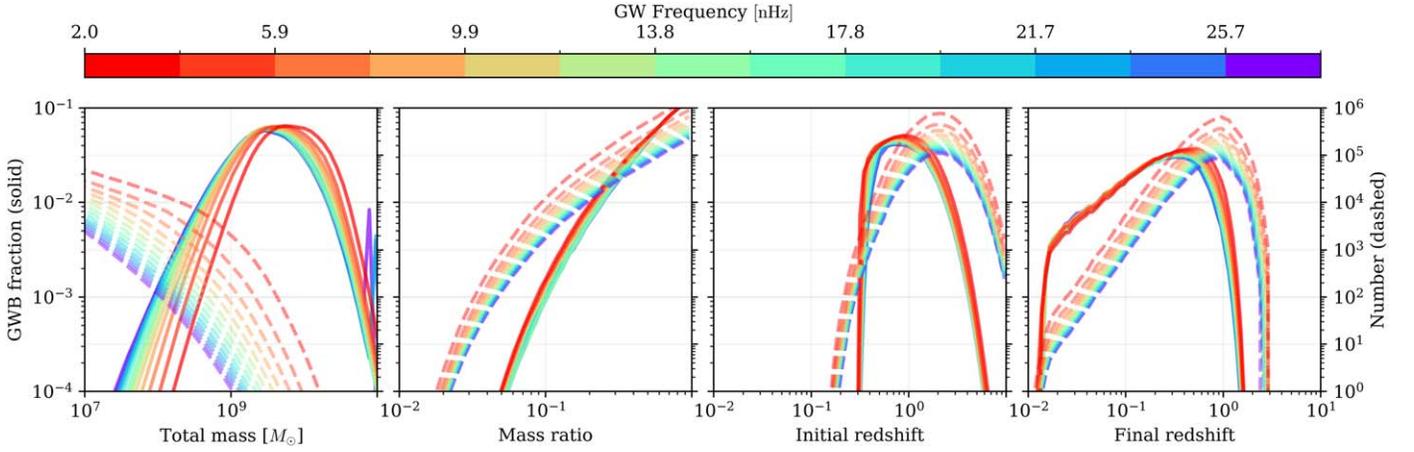


Figure 12. Comparison of binary parameters by their fractional contribution to the GWB (h_c^2 ; solid) and their total number of binaries (dashed). Line colors indicate the GW frequency at which the binaries emit. The third panel (z_{init}) corresponds to the “initial” galaxy merger redshift, while the fourth panel (z_{final}) is the redshift at which the binary is emitting GWs. This figure demonstrates that the GW signal is produced by a relatively small and highly biased subsample of a much larger population.

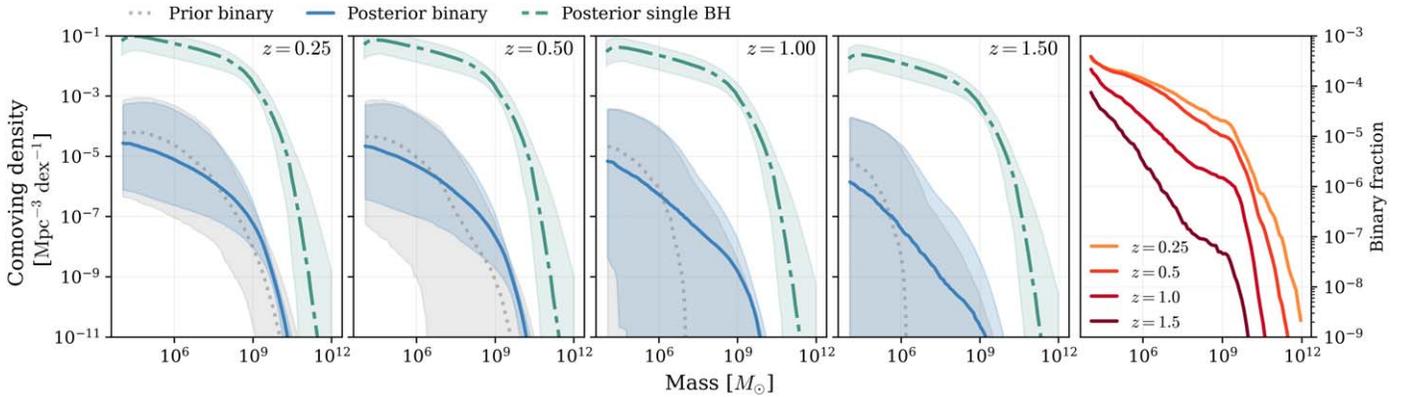


Figure 13. The first four panels show comoving number density per logarithmic interval of mass. Densities are shown for binary black holes in the NANOGrav frequency band at the redshift indicated in the upper right corner for the priors (gray dotted line) and posteriors (blue solid line) derived from our *Phenom* population models. Densities are also shown for the total SMBH population (teal dashed-dotted line). Shaded regions show the 68% distributions. The last panel shows the binary fraction as a function of mass for the four selected redshifts. The implied mass functions of our posteriors prefer a relatively high density of black holes larger than $M = 10^9 M_\odot$, and our posteriors are overall more confined than our priors at lower redshifts.

future NANOGrav 20 yr data set would be far more constraining than the 15 yr data. In that case, the authors found that subtle differences in the degree of environmental coupling could be distinguished based on spectra with nearly identical reference amplitudes but differing low-frequency spectral shapes.

While we are unable to definitively attribute the inferred GWB signal to SMBH binaries at the current signal-to-noise ratio, we show that all of the signal’s features are consistent with binaries. Nonetheless, many other possible origins of the GWB have been proposed, as detailed in [NG15newphys](#). It is worth emphasizing that SMBH binaries must necessarily form throughout the Universe as a natural product of galaxy mergers. If the inferred GWB is not dominantly produced by SMBH binaries, the lack of GW signal from inspiraling SMBH binaries must somehow be accounted for. One possibility is that SMBH binaries usually stall outside of the GW-driven regime of inspiral, which could occur if gas- and stellar-driven processes are insufficient to bring binaries to the GW-dominated regime. If the so-called “final-parsec problem” (e.g., [Begelman et al. 1980](#)) indeed lengthens most inspiral timescales to a Hubble time or longer, the resulting GWB from binaries could be attenuated to amplitudes well below the

inferred 15 yr signal. Even in this pessimistic case, multiple studies have suggested that triple SMBH interactions would still produce a detectable GWB signal ([Volonteri et al. 2003](#); [Hoffman & Loeb 2007](#); [Bonetti et al. 2018b](#); [Ryu et al. 2018](#)). In either case, this would also imply the existence of a large population of stalled SMBH binaries in the local Universe.

Additional data are required to resolve the origin of the GWB. One of the strongest distinguishing features between different source models is the significantly higher degree of anisotropy for binaries as opposed to new physical processes ([NG15newphys](#)). From the 15 yr data set, the first limits on anisotropy have now been placed ([Agazie et al. 2023e](#)). While the limits are still consistent with astrophysical expectations for binary populations (e.g., [Sato-Polito & Kamionkowski 2023](#)), they will become significantly more constraining over time ([Ali-Haïmoud et al. 2021](#); [Pol et al. 2022](#)).

Eventually, individual continuous-wave GW sources will also become distinguishable above the GWB, if it is indeed produced by binaries. Different models have produced a variety of expectations for the plausibility of continuous-wave source detection in the near future ([Sesana et al. 2009](#); [Rosado et al. 2015](#); [Mingarelli et al. 2017](#); [Kelley et al. 2018](#); [Bécsy et al. 2022](#)). A search for continuous-wave sources has yielded

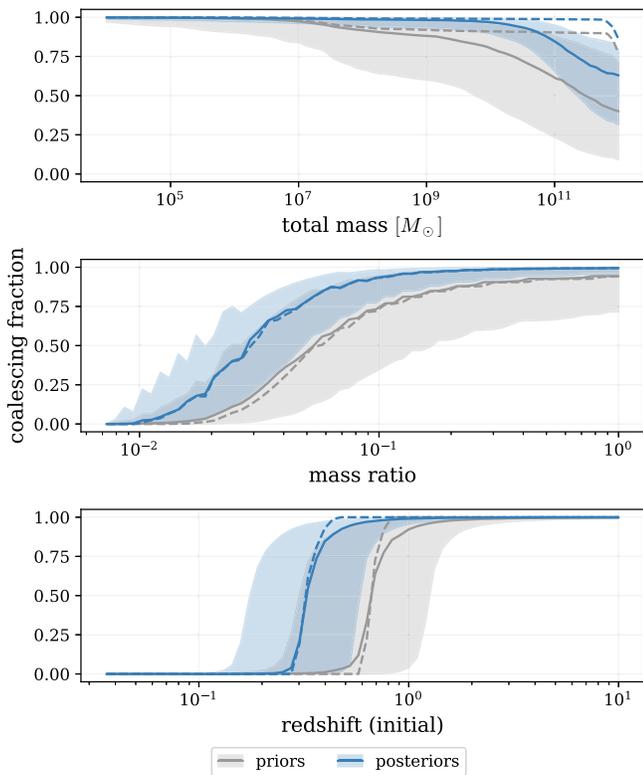


Figure 14. Fraction of binary systems that reach the PTA band before redshift zero in our simulated populations. The top panel shows the total binary mass M , the middle panel shows the binary mass ratio q , and the bottom panel shows the initial binary redshift (defined as the galaxy merger redshift) z_{init} . Solid lines show median values over 1000 samples from the priors (gray) and posteriors (blue) after fitting to the 15 yr data. The shaded regions correspond to the 50% interquartile range. The dashed lines correspond to binary subsets, with $q > 0.2$ in the top panel, $M > 10^8 M_{\odot}$ in the middle panel, and both cuts in the bottom panel. Fits to the 15 yr data strongly favor shorter binary lifetimes, allowing larger fractions of binaries to coalesce. The nearly step-function behavior in the bottom panel marks the redshift at which the look-back time matches the sample’s binary-lifetime parameter.

improved upper limits on their occurrence rates in both the 12.5 (Arzoumanian et al. 2023) and 15 yr data (Agazie et al. 2023d). A continuous-wave detection would present the exciting possibility of multimessenger detections: GWs from a single SMBH binary for which an electromagnetic counterpart could be identified. Such a multimessenger source would provide a wealth of information about the origin of low-frequency GWs, the astrophysical environment of SMBH binaries, and SMBH accretion processes (Kelley et al. 2019a).

In the next decade, the Laser Interferometer Space Antenna (LISA) will begin operation in the \sim millihertz band, sensitive to the merger of SMBH binaries with masses in a range between $\sim 10^4$ and $10^8 M_{\odot}$ out to $z \sim 10$ (Amaro-Seoane et al. 2023). While LISA promises to reveal the elusive formation mechanism of massive black hole seeds in the early Universe, this range of binary masses is more poorly constrained and far more challenging to model than the higher-mass PTA binaries. The approaches and analyses developed for studying PTA GW sources will be crucial for paving the way for LISA science. The identification of LISA electromagnetic counterparts will also be much more difficult, further motivating the development of techniques for PTA sources and generalizing them to signals at lower masses.

If the GWB signal is indeed produced by astrophysical binaries, it will be the first proof that SMBH binaries do indeed

form, evolve to subparsec separations, and eventually coalesce. These systems will join the new landscape of multimessenger GW astrophysics, offering the opportunity to study the most extreme and energetic environments in the Universe and probe the closely coupled coevolution of galaxies and their nuclear engines. If, instead, the GWB has a different cosmological origin, it may provide answers to the most outstanding questions in fundamental physics that challenge the standard model and Λ CDM. In either case, PTAs have cracked open the era of low-frequency GW astronomy.

Acknowledgments

Author contributions. An alphabetical-order author list was used for this paper in recognition of the fact that a large, decade-timescale project such as NANOGrav is necessarily the result of the work of many people. All authors contributed to the activities of the NANOGrav collaboration leading to the work presented here and reviewed the manuscript, text, and figures prior to the paper’s submission. Additional specific contributions to this paper are as follows. L.Z.K. proposed and coordinated the project. Development of the `holodeck` population modeling framework was led by L.Z.K., with contributions from A.C.-C., D.W., E.C.G., J.M.W., K.G., M.S.S., and S.C. Libraries were generated by J.S., K.G., L.B., L.Z.K., and N.G.-D. Development of the Gaussian process modules was led by J.S., with contributions from D.W. and J.M.W. Development of the fitting infrastructure was led by W.G.L., with contributions from J.S. Gaussian process training was performed by D.W. and J.S., with diagnostics performed by J.M.W. The 15 yr GWB data products were provided by the NANOGrav Detection Working Group. The analysis was performed by J.S., K.G., and L.Z.K., who produced most of the figures, and by L.B. Additional analysis code and interpretation were provided by D.J.D. and E.C.G. The collection of literature predictions and production of the literature review figure was performed by M.C. The paper was written by J.C.R., J.S., K.G., L.B., L.Z.K., and M.C.

A.Mi., B.B., D.J.D., J.S.H., M.T.L., M.V., S.J.V., S.R.T., T. J.W.L., and T.D. provided feedback on the analysis and manuscript.

The NANOGrav collaboration receives support from National Science Foundation (NSF) Physics Frontiers Center award Nos. 1430284 and 2020265, the Gordon and Betty Moore Foundation, NSF AccelNet award No. 2114721, an NSERC Discovery Grant, and CIFAR. NANOGrav is part of the International Pulsar Timing Array (IPTA).

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.

This work was conducted in part using the resources of the Advanced Computing Center for Research and Education (ACCRES) at Vanderbilt University, Nashville, TN. This work received computational support from UCF’s Advanced Research Computing cluster, operated by the University of Central Florida. This research used the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by

the UC Berkeley Chancellor, Vice Chancellor for Research, and Chief Information Officer). This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor. This work was also conducted in part using the Thorny Flat HPC Cluster at West Virginia University (WVU), which is funded in part by National Science Foundation (NSF) Major Research Instrumentation Program (MRI) Award No. 1726534 and West Virginia University. This work utilized the Alpine high performance computing resource at the University of Colorado Boulder. Alpine is jointly funded by the University of Colorado Boulder, the University of Colorado Anschutz, Colorado State University, and the National Science Foundation (award 2201538). This work utilized the Blanca condo computing resource at the University of Colorado Boulder. Blanca is jointly funded by computing users and the University of Colorado Boulder.

L.B. acknowledges support from the National Science Foundation under award AST-1909933 and the Research Corporation for Science Advancement under Cottrell Scholar Award No. 27553. P.R.B. is supported by the Science and Technology Facilities Council, 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., D.D., N.La., X.S., J.P.S., and J.T. is partly supported by the George and Hannah Bolinger Memorial Fund in the College of Science at Oregon State University. M.C., P. P., J.C.R., and S.R.T. acknowledge support from NSF AST-2007993. M.C. and N.S.P. were supported by the Vanderbilt Initiative in Data Intensive Astrophysics (VIDA) Fellowship. K.Ch., A.D.J., and M.V. acknowledge support from the Caltech and Jet Propulsion Laboratory President’s and Director’s Research and Development Fund. K.Ch. and A.D. J. acknowledge support from the Sloan Foundation. Support for this work was provided by the NSF through the Grote Reber Fellowship Program administered by Associated Universities, Inc./National Radio Astronomy Observatory. Support for H.T. C. is provided by NASA through NASA Hubble Fellowship Program grant No. HST-HF2-51453.001 awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. K.Cr. is supported by a UBC Four Year Fellowship (6456). M.E.D. acknowledges support from the Naval Research Laboratory by NASA under contract S-15633Y. T.D. and M.T.L. are supported by NSF Astronomy and Astrophysics Grant (AAG) award No. 2009468. E.C.F. is supported by NASA under award No. 80GSFC21M0002. G.E.F., S.C.S., and S.J.V. are supported by NSF award PHY-2011772. K.A.G. and S.R.T. acknowledge support from NSF CAREER award No. 2146016. The Flatiron Institute is supported by the Simons Foundation. S.H. is supported by the National Science Foundation Graduate Research Fellowship under grant No. DGE-1745301. N.La. acknowledges the support from the Larry W. Martin and Joyce B. O’Neill Endowed Fellowship in the College of Science at Oregon State University. 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). D.R.L. and M.A.Mc. are supported by NSF No. 1458952. M.A.Mc. is supported by NSF No. 2009425. C.M.F.M. was supported in part by the National Science Foundation 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. P.N. acknowledges support from the BHI, funded by grants from the John Templeton Foundation and the Gordon and Betty Moore Foundation. 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. 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. J.S. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-2202388 and acknowledges previous support by the NSF under award 1847938. C.U. acknowledges support from BGU (Kreitman fellowship) and the Council for Higher Education and Israel Academy of Sciences and Humanities (Excellence fellowship). C.A.W. acknowledges support from CIERA, the Adler Planetarium, and the Brinson Foundation through a CIERA-Adler postdoctoral fellowship. O.Y. is supported by the National Science Foundation Graduate Research Fellowship under grant No. DGE-2139292. D.J.D. received funding from the European Union’s Horizon 2020 research and innovation program under Marie Skłodowska-Curie grant agreement No. 101029157 and the Danish Independent Research Fund through Sapere Aude Starting Grant No. 121587.

Facilities: Arecibo, GBT, VLA.

Software: acor, astropy (Astropy Collaboration et al. 2018), ceflyl (Lamb et al. 2023), chainconsumer (Hinton 2016), cython (Behnel et al. 2011), emcee (Foreman-Mackey et al. 2013), enterprise (Ellis et al. 2020), enterprise_extensions (Taylor et al. 2021), George (Ambikasaran et al. 2015b), holodeck (L. Z. Kelley et al. 2023, in preparation), jupyter (Kluyver et al. 2016), kalepy (Kelley 2021b), matplotlib (Hunter 2007), numba (Lam et al. 2015), numpy (van der Walt et al. 2011), ptmcmc (Ellis & van Haasteren 2017), scipy (Virtanen et al. 2020).

Appendix A GWB Predictions in the Literature

In Figure A1 and Table A1, we summarize the model predictions of GWB amplitudes from the literature. For each reference amplitude, predictions are cited at frequencies of $(1 \text{ yr})^{-1}$ and $(10 \text{ yr})^{-1}$. While it is most common to reference GWB amplitudes at $f = 1 \text{ yr}^{-1}$, NANOGrav constraints are derived primarily at much lower frequencies. Small deviations in power-law indices at low frequencies can lead to large changes in the corresponding A_{yr} amplitudes. The amplitudes at $f = (10 \text{ yr})^{-1}$ are much more representative of current PTA constraints. Figure A1 also shows the posterior distributions for the inferred GWB amplitudes obtained from the 15 yr NANOGrav data when assuming a power-law model; results are shown for both the *HD-w/MP+DP+CURN* and *HD-DMGP* models. Numerous literature model predictions overlap with the inferred GWB amplitudes from the 15 yr data, especially when comparing the amplitude at $f = (10 \text{ yr})^{-1}$. Even so, the 15 yr results lie at the higher-amplitude end of the predicted ranges.

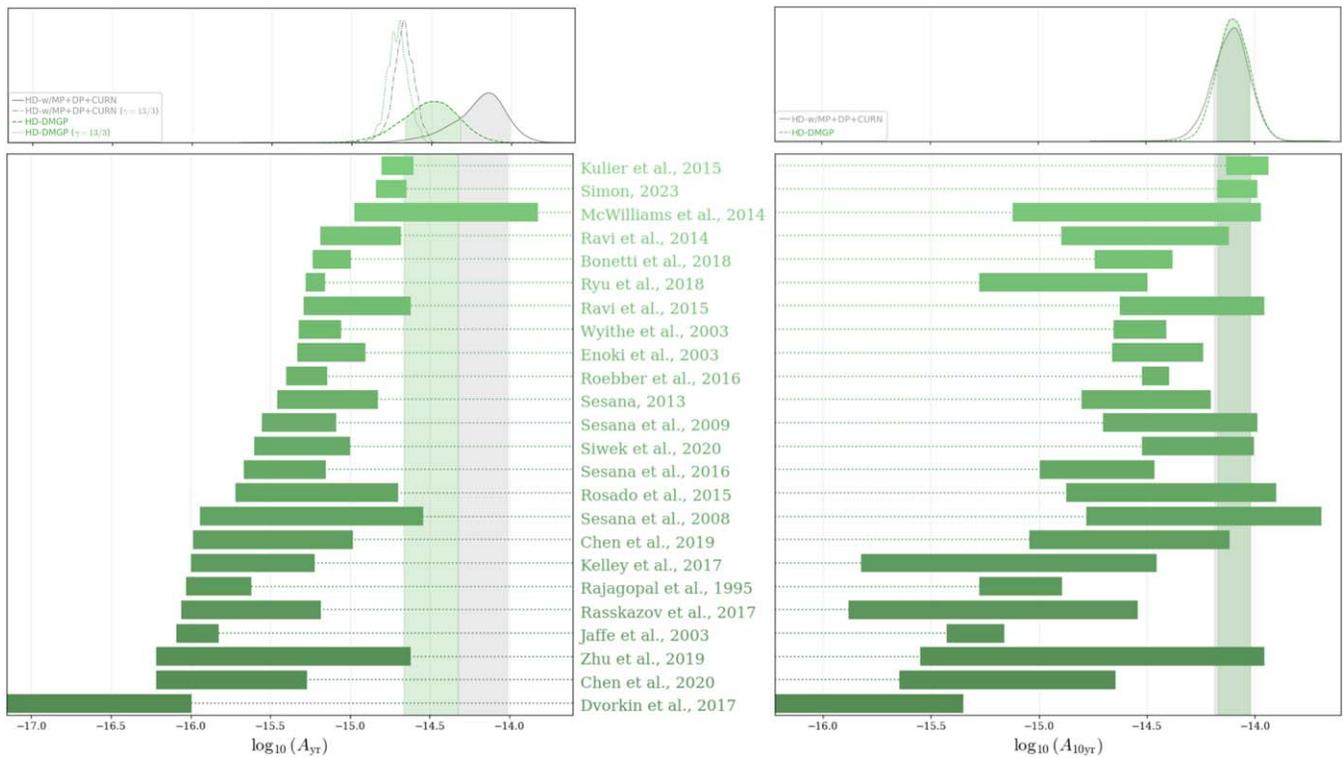


Figure A1. Literature predictions for the amplitudes A_{yr} and $A_{10 \text{ yr}}$ of the GWB at frequencies of $(1 \text{ yr})^{-1}$ and $(10 \text{ yr})^{-1}$, respectively, compared to the NANOGrav 15 yr results. While it is most common to reference GWB amplitudes at $f = 1 \text{ yr}^{-1}$, NANOGrav constraints are primarily derived at lower frequencies; thus, the $f = (10 \text{ yr})^{-1}$ values are much more representative of current PTA constraints. The green horizontal bars indicate the 16th–84th percentile uncertainty regions for each prediction; these model predictions are also listed in Table A1. The amplitude distributions in the upper panels correspond to the posterior probability distributions of GWB amplitude for power-law models fit to the GWB free-spectrum posteriors (*HD-w/MP+DP+CURN*, gray solid curves, and *HD-DMGP*, green dashed curves; see Figure 5 in [NG15gwb](#)). The $\gamma = 13/3$ slice is also shown for the A_{yr} values with dashed–dotted and dotted lines for the *HD-w/MP+DP+CURN* and *HD-DMGP* values, respectively, since those are the most directly comparable to many of the models included here. The green and gray shaded regions are the corresponding 68% credible intervals. The 15 yr NANOGrav results are within the bounds of some model predictions but require GWB amplitudes at the higher end of the predicted ranges.

Table A1
Literature Predictions for the Amplitudes A_{yr} and $A_{10\text{yr}}$ of the GWB at Frequencies of $(1 \text{ yr})^{-1}$ and $(10 \text{ yr})^{-1}$, Respectively

Model (1)	A_{yr} 16th (2)	A_{yr} 84th (3)	$A_{10\text{yr}}$ 16th (4)	$A_{10\text{yr}}$ 84th (5)
Rajagopal & Romani (1995)	9.32E-17	2.41E-16	5.31E-16	1.28E-15
Jaffe & Backer (2003)	8.10E-17	1.50E-16	3.75E-16	6.93E-16
Wyithe & Loeb (2003)	4.77E-16	8.84E-16	2.22E-15	3.89E-15
Enoki et al. (2004)	4.70E-16	1.25E-15	2.18E-15	5.77E-15
Sesana et al. (2008)	1.15E-16	2.88E-15	1.66E-15	2.04E-14
Sesana et al. (2009)	2.79E-16	8.21E-16	1.98E-15	1.03E-14
Sesana (2013b)	3.50E-16	1.50E-15	1.58E-15	6.26E-15
McWilliams et al. (2014)	1.07E-15	1.51E-14	7.58E-16	1.07E-14
Ravi et al. (2014)	6.51E-16	2.10E-15	1.27E-15	7.62E-15
Kulier et al. (2015)	1.58E-15	2.51E-15	7.36E-15	1.16E-14
Ravi et al. (2015)	5.10E-16	2.40E-15	2.37E-15	1.11E-14
Rosado et al. (2015)	1.91E-16	2.01E-15	1.34E-15	1.26E-14
Roebber et al. (2016)	4.00E-16	7.23E-16	3.00E-15	4.00E-15
Sesana et al. (2016)	2.15E-16	7.08E-16	1.01E-15	3.43E-15
Rasskazov & Merritt (2017)	8.74E-17	6.57E-16	1.32E-16	2.87E-15
Dvorkin & Barausse (2017)	8.74E-17	6.57E-16	1.32E-16	2.87E-15
Kelley et al. (2017b)	1.00E-16	6.00E-16	1.50E-16	3.50E-15
Ryu et al. (2018)	5.30E-16	7.00E-16	5.30E-16	3.20E-15
Bonetti et al. (2018b)	5.83E-16	1.01E-15	1.81E-15	4.18E-15
Zhu et al. (2019)	6.10E-17	2.40E-15	2.83E-16	1.11E-14
Chen et al. (2019)	1.04E-16	1.05E-15	9.02E-16	7.63E-15
Chen et al. (2020)	6.10E-17	5.40E-16	2.26E-16	2.27E-15
Siwek et al. (2020)	2.50E-16	1.00E-15	3.00E-15	9.94E-15
Simon (2023)	1.46E-15	2.26E-15	6.67E-15	1.03E-14

Note. Column (1) gives the literature reference, columns (2) and (3) give the 16th and 84th percentiles for the uncertainty region of the corresponding prediction of A_{yr} , and columns (4) and (5) give the 16th and 84th percentiles for $A_{10\text{yr}}$ predictions. Figure A1 provides a visual comparison of these predictions to the NANOGrav 15 yr results.

Appendix B Semianalytic Model Parameterizations

In Section 3.2, we presented the semianalytic models and underlying equations used in this paper. In Table B1, we detail

all of the model parameters, showing the fiducial values for fixed parameters, as well as the prior distributions for parameters that are varied when fitting the models to the 15 yr data. These models are summarized in Table B2.

Table B1
Astrophysical Parameters of our Semianalytic Population Models

Model Component	Symbol	Fiducial Value	Uniform Priors	Astrophysical Priors
GSMF ^a (Ψ)	ψ_0	...	$\mathcal{U}(-3.5, -1.5)$	$\mathcal{N}(-2.56, 0.4)$
	ψ_z	-0.60
	$m_{\psi,0}$...	$\mathcal{U}(10.5, 12.5)$	$\mathcal{N}(10.9, 0.4)$
	$m_{\psi,z}$	+0.11
	$\alpha_{\psi,0}$	-1.21	...	$\mathcal{N}(-1.2, 0.2)$
	$\alpha_{\psi,z}$	-0.03
GPF ^b (P)	P_0	+0.033
	$\alpha_{p,0}$	0.0
	$\alpha_{p,z}$	0.0
	$\beta_{p,0}$	+1.0	...	$\mathcal{N}(0.8, 0.4)$
	$\beta_{p,z}$	0.0
	$\gamma_{p,0}$	0.0	...	$\mathcal{N}(0.5, 0.3)$
	$\gamma_{p,z}$	0.0
GMT ^c ($T_{\text{gal-gal}}$)	T_0	+0.5 Gyr	...	$\mathcal{U}(0.2, 5.0)$ Gyr
	$\alpha_{t,0}$	0.0
	$\alpha_{t,z}$	0.0
	$\beta_{t,0}$	-0.5	...	$\mathcal{U}(-2.0, 0.0)$
	$\beta_{t,z}$	0.0
	$\gamma_{t,0}$	-1.0
	$\gamma_{t,z}$	0.0
$M_{\text{BH}}-M_{\text{bulge}}$ ^d (M_{BH})	μ	...	$\mathcal{U}(7.6, 9.0)$	$\mathcal{N}(8.6, 0.2)$
	α_μ	+1.10	...	$\mathcal{N}(1.2, 0.2)$
	ϵ_μ	...	$\mathcal{U}(0.0, 0.9)$ dex	$\mathcal{N}(0.32, 0.15)$ dex
	$f_{*,\text{bulge}}$	+0.615
Phenom ($\frac{da}{dt}$)	τ_f	...	$\mathcal{U}(0.1, 11.0)$ Gyr	$\mathcal{U}(0.1, 11.0)$ Gyr
	a_c	+10 ² pc
	a_{init}	+10 ³ pc
	ν_{inner}	...	$\mathcal{U}(-1.5, 0.0)$	$\mathcal{U}(-1.5, +0.5)$
	ν_{outer}	+2.5

Notes. Units are denoted where relevant; all other parameters are defined to be dimensionless. For libraries, we denote uniform distributions with $\mathcal{U}(\text{min}, \text{max})$ and normal distributions with $\mathcal{N}(\text{mean}, \text{std. dev.})$.

^a The fiducial GSMF values are based on Chen et al. (2019), while the ‘‘astrophysical library’’ parameters are based on fits to the data from Tomczak et al. (2014).

^b The GPF parameters are based on a comparison of Conselice et al. (2003), Bluck et al. (2012), Mundy et al. (2017), and Duncan et al. (2019).

^c The GMT parameters are based on a comparison of Conselice et al. (2008), Boylan-Kolchin et al. (2008), Conselice (2009), and Snyder et al. (2017).

^d The $M_{\text{BH}}-M_{\text{bulge}}$ parameters are based on Gültekin et al. (2009b), Kormendy & Ho (2013), and McConnell & Ma (2013), with bulge fractions based on Lang et al. (2014) and Bluck et al. (2014).

Table B2
Summary of Semianalytic SMBH Binary Population Models Used in This Work

Model Name	Parameters Varied (by Model Component)					Priors
	GSMF	GPF	GMT	$M_{\text{BH}}-M_{\text{bulge}}$	Phenom	
<i>Phenom+Uniform</i>	$\psi_0, m_{\psi,0}$	μ, ϵ_μ	$\tau_f, \nu_{\text{inner}}$	Uniform
<i>Phenom+Astro</i>	$\psi_0, m_{\psi,0}$	μ, ϵ_μ	$\tau_f, \nu_{\text{inner}}$	Astrophysical
<i>GWOnly+Uniform</i>	$\psi_0, m_{\psi,0}$	μ, ϵ_μ	...	Uniform
<i>Phenom-Ext+Astro</i>	$\psi_0, m_{\psi,0}, \alpha_{\psi,0}$	$\beta_{p,0}, \gamma_{p,0}$	$T_0, \beta_{t,0}$	$\mu, \alpha_\mu, \epsilon_\mu$	$\tau_f, \nu_{\text{inner}}$	Astrophysical
<i>GWOnly-Ext+Astro</i>	$\psi_0, m_{\psi,0}, \alpha_{\psi,0}$	$\beta_{p,0}, \gamma_{p,0}$	$T_0, \beta_{t,0}$	$\mu, \alpha_\mu, \epsilon_\mu$...	Astrophysical

Note. Model parameters are defined in Section 3.2, and their fiducial values and assumed prior distributions are given in Table B1. The first model (*Phenom+Uniform*), indicated in bold, is what we refer to as our fiducial model. Throughout the text, when the assumed priors can be omitted from the model name without loss of clarity, we simply refer to the models as *Phenom*, *GWOnly*, *Phenom-Ext*, and *GWOnly-Ext*. The latter two (also referred to as the “extended models”) are discussed in Appendix C.

Appendix C Higher-dimensional Parameter Spaces

We have shown that in our fiducial, six-dimensional *Phenom* parameter space, a wide range of semianalytic model parameters are consistent with current measurements of the GWB. In this library, a large number of additional parameters are held fixed to astrophysically motivated values. When fitting to the 15 yr GWB data, uniform priors on the included parameters are typically used. We have compared these results to fits of the same parameter space but adopting more informed priors based on the astronomical literature.

Currently, it is not feasible to run MCMC fits using GPs that have been trained to significantly larger parameter spaces. However, we have generated higher-dimensional libraries and directly evaluated them against the 15 yr GWB spectra at the

library grid points themselves. In this way, we can weight the input parameters by the resulting likelihoods to obtain posteriors without using MCMC to dynamically explore the domain. This approach allows us to examine the effects of freeing additional parameters.

Figure C1 compares the parameter posteriors for our fiducial library against the larger parameter spaces, including a 12-dimensional phenomenological version (*Phenom-Ext*; orange dashed) and 10-dimensional GW-only version (*GWOnly-Ext*; purple dashed). Blue solid lines show the standard *Phenom* fits with uniform priors, while orange solid lines show the same six-dimensional library but fitting with priors taken from the *Phenom-Ext* distributions.

Posteriors from the six-dimensional phenomenological models are entirely consistent with the 10- and 12-dimensional

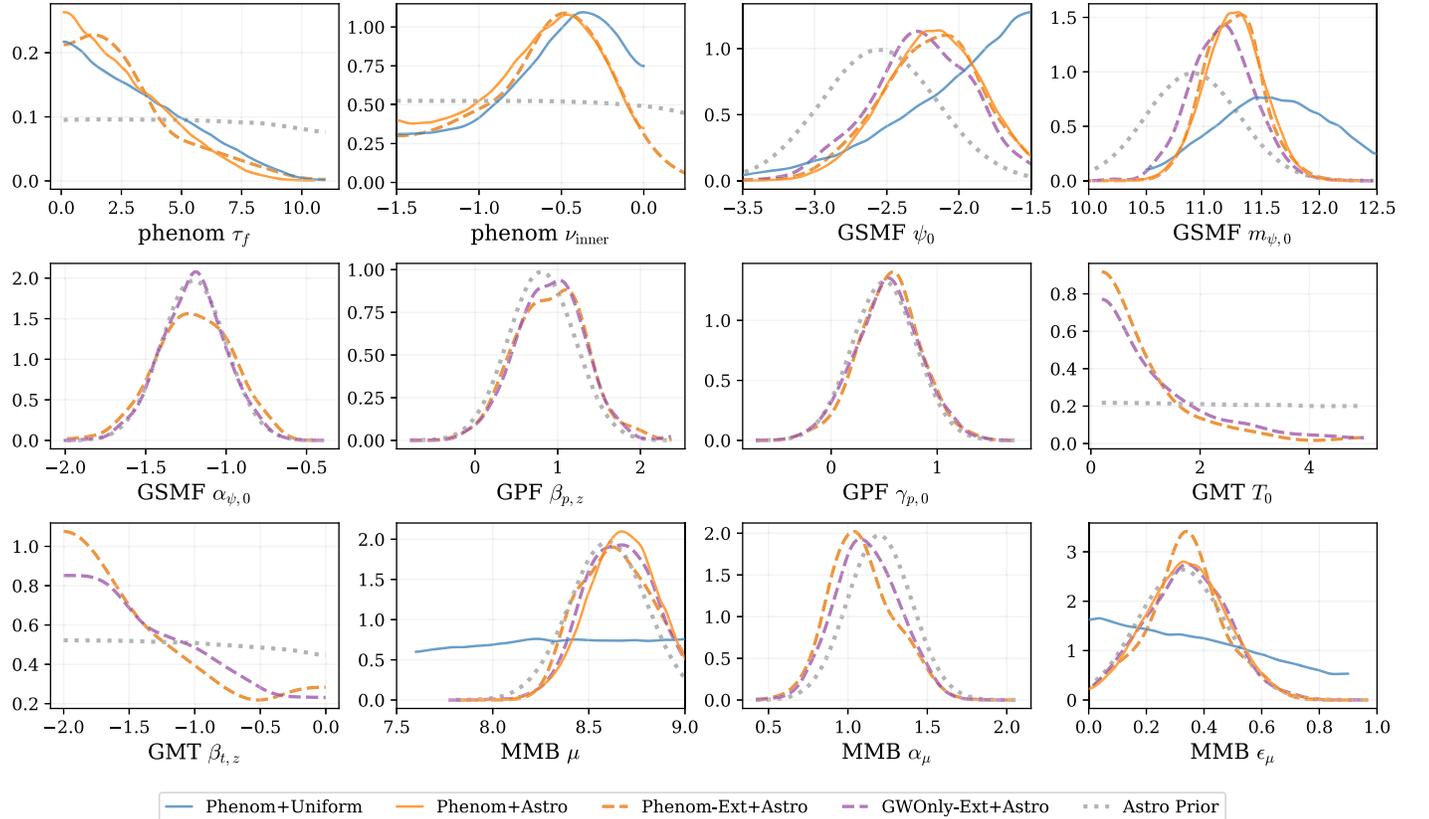


Figure C1. Semianalytic model posteriors comparing our fiducial six-dimensional *Phenom* model (solid) to our much larger *Phenom-Ext* parameter space (dashed). The standard version of the *Phenom* library with uniform priors (blue) is also compared to a version that is fit against the 15 yr GWB spectra using the same astrophysically motivated priors from the *Phenom+Astro* distributions (orange). The priors are also shown for the astrophysically motivated case (gray dotted).

Phenom-Ext libraries when using the same priors. Setting the additional parameters to fixed values does not bias the resulting measurements, nor does it lead to underestimating the width of the posterior distributions. This is likely the case because the current 15 yr NANOGrav spectral measurements include large uncertainties. As the data improve, it will become more important to fully explore the parameter space.

As previously discussed, for many parameters, the shape of the priors does significantly impact the recovered posteriors. This is particularly noticeable in the GSMF (ψ_0 and $m_{\psi,0}$) and $M_{\text{BH}}-M_{\text{bulge}}$ (μ and ϵ_{μ}) parameters, where there is a significant degeneracy between parameters that broadly changes the amplitude of the GWB spectrum.

ORCID iDs

Gabriella Agazie <https://orcid.org/0000-0001-5134-3925>

Akash Anumrapudi <https://orcid.org/0000-0002-8935-9882>

Anne M. Archibald <https://orcid.org/0000-0003-0638-3340>

Paul T. Baker <https://orcid.org/0000-0003-2745-753X>

Bence Bécsy <https://orcid.org/0000-0003-0909-5563>

Laura Blecha <https://orcid.org/0000-0002-2183-1087>

Alexander Bonilla <https://orcid.org/0000-0002-7001-0728>

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>

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>

Katerina Chatziioannou <https://orcid.org/0000-0002-5833-413X>

Siyuan Chen <https://orcid.org/0000-0002-3118-5963>

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>

Curt J. Cutler <https://orcid.org/0000-0002-2080-1468>

Daniel J. D’Orazio <https://orcid.org/0000-0002-1271-6247>

Megan E. DeCesar <https://orcid.org/0000-0002-2185-1790>

Paul B. Demorest <https://orcid.org/0000-0002-6664-965X>

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 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>

Joseph Glaser <https://orcid.org/0000-0003-4090-9780>

Deborah C. Good <https://orcid.org/0000-0003-1884-348X>

Kayhan Gültekin <https://orcid.org/0000-0002-1146-0198>

Jeffrey S. Hazboun <https://orcid.org/0000-0003-2742-3321>

Sophie Hourihane <https://orcid.org/0000-0002-9152-0719>

Ross J. Jennings <https://orcid.org/0000-0003-1082-2342>

Aaron Johnson <https://orcid.org/0000-0002-7445-8423>

Megan L. Jones <https://orcid.org/0000-0001-6607-3710>

Andrew R. Kaiser <https://orcid.org/0000-0002-3654-980X>

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>

Natalia Lewandowska <https://orcid.org/0000-0003-0771-6581>

Tyson B. Littenberg <https://orcid.org/0000-0002-9574-578X>

Tingting Liu <https://orcid.org/0000-0001-5766-4287>

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>

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>

Priyamvada Natarajan <https://orcid.org/0000-0002-5554-8896>

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>

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>

Levi Schult <https://orcid.org/0000-0001-6425-7807>

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>

Magdalena S. Siwek <https://orcid.org/0000-0002-1530-9778>

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>
 Jerry P. Sun  <https://orcid.org/0000-0002-7933-493X>
 Abhimanyu Susobhanan  <https://orcid.org/0000-0002-2820-0931>
 Joseph K. Swiggum  <https://orcid.org/0000-0002-1075-3837>
 Stephen R. Taylor  <https://orcid.org/0000-0001-8217-1599>
 Jacob E. Turner  <https://orcid.org/0000-0002-2451-7288>
 Caner Unal  <https://orcid.org/0000-0001-8800-0192>
 Michele Vallisneri  <https://orcid.org/0000-0002-4162-0033>
 Sarah J. Vigeland  <https://orcid.org/0000-0003-4700-9072>
 Jeremy M. Wächter  <https://orcid.org/0000-0002-1070-2431>
 Haley M. Wahl  <https://orcid.org/0000-0001-9678-0299>
 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

- Afzal, A., Agazie, G., Anumarlapudi, A., et al. 2023, *ApJL*, 951, L11
 Agazie, G., Anumarlapud, A., Archibald, A. M., et al. 2023a, *ApJL*, 951, L8
 Agazie, G., Alam, M. F., Anumarlapudi, A., et al. 2023b, *ApJL*, 951, L9
 Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023c, *ApJL*, 951, L10
 Agazie, G., Baker, P. T., Blecha, L., et al. 2023d, arXiv:2301.03608
 Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023e, arXiv:2306.16221
 Aigrain, S., & Foreman-Mackey, D. 2022, arXiv:2209.08940
 Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021, *ApJS*, 252, 4
 Ali-Haïmoud, Y., Smith, T. L., & Mingarelli, C. M. F. 2021, *PhRvD*, 103, 042009
 Amaro-Seoane, P., et al. 2023, *LRR*, 26, 2
 Ambikasaran, S., Foreman-Mackey, D., Greengard, L., Hogg, D. W., & O’Neil, M. 2015a, *ITPAM*, 38, 252
 Ambikasaran, S., Foreman-Mackey, D., Greengard, L., Hogg, D. W., & O’Neil, M. 2015b, *ITPAM*, 38, 252
 Antoniadis, J., Arzoumanian, Z., Babak, S., et al. 2022, *MNRAS*, 510, 4873
 Antonini, F., & Merritt, D. 2012, *ApJ*, 745, 83
 Arzoumanian, Z., Baker, P. T., Blecha, L., et al. 2023, arXiv:2301.03608
 Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, 905, L34
 Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018, *ApJ*, 859, 47
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
 Bansal, K., Taylor, G. B., Peck, A. B., Zavala, R. T., & Romani, R. W. 2017, *ApJ*, 843, 14
 Barnes, J. E., & Hernquist, L. 1992, *ARA&A*, 30, 705
 Bécsy, B., Cornish, N. J., & Kelley, L. Z. 2022, *ApJ*, 941, 119
 Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, *Natur*, 287, 307
 Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, 13, 31
 Berczik, P., Merritt, D., Spurzem, R., & Bischof, H.-P. 2006, *ApJL*, 642, L21
 Bluck, A. F. L., Conselice, C. J., Buitrago, F., et al. 2012, *ApJ*, 747, 34
 Bluck, A. F. L., Mendel, J. T., Ellison, S. L., et al. 2014, *MNRAS*, 441, 599
 Bogdanović, T., Miller, M. C., & Blecha, L. 2022, *LRR*, 25, 3
 Bonetti, M., Haardt, F., Sesana, A., & Barausse, E. 2016, *MNRAS*, 461, 4419
 Bonetti, M., Haardt, F., Sesana, A., & Barausse, E. 2018a, *MNRAS*, 477, 3910
 Bonetti, M., Sesana, A., Barausse, E., & Haardt, F. 2018b, *MNRAS*, 477, 2599
 Bowen, D. B., Mewes, V., Campanelli, M., et al. 2018, *ApJL*, 853, L17
 Boylan-Kolchin, M., Ma, C.-P., & Quataert, E. 2008, *MNRAS*, 383, 93
 Breiding, P., Burke-Spolaor, S., Eracleous, M., et al. 2021, *ApJ*, 914, 37
 Burke-Spolaor, S. 2011, *MNRAS*, 410, 2113
 Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, *A&ARv*, 27, 5
 Casey-Clyde, J. A., Mingarelli, C. M. F., Greene, J. E., et al. 2022, *ApJ*, 924, 93
 Chandrasekhar, S. 1943, *ApJ*, 97, 255
 Charisi, M., Bartos, I., Haiman, Z., et al. 2016, *MNRAS*, 463, 2145
 Charisi, M., Haiman, Z., Schiminovich, D., & D’Orazio, D. J. 2018, *MNRAS*, 476, 4617
 Chen, S., Caballero, R. N., Guo, Y. J., et al. 2021, *MNRAS*, 508, 4970
 Chen, S., Sesana, A., & Conselice, C. J. 2019, *MNRAS*, 488, 401
 Chen, S., Sesana, A., & Del Pozzo, W. 2017, *MNRAS*, 470, 1738
 Chen, Y., Yu, Q., & Lu, Y. 2020, *ApJ*, 897, 86
 Chen, Y.-C., Hwang, H.-C., Shen, Y., et al. 2022, *ApJ*, 925, 162
 Comerford, J. M., Gerke, B. F., Newman, J. A., et al. 2009, *ApJ*, 698, 956
 Comerford, J. M., & Greene, J. E. 2014, *ApJ*, 789, 112
 Conselice, C. J. 2009, *MNRAS*, 399, L16
 Conselice, C. J., Bershady, M. A., Dickinson, M., & Papovich, C. 2003, *AJ*, 126, 1183
 Conselice, C. J., Rajgor, S., & Myers, R. 2008, *MNRAS*, 386, 909
 Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *MNRAS*, 365, 11
 Curylo, M., & Bulik, T. 2022, *A&A*, 660, A68
 De Rosa, A., Vignali, C., Bogdanović, T., et al. 2019, *NewAR*, 86, 101525
 Di Matteo, T., Springel, V., & Hernquist, L. 2005, *Natur*, 433, 604
 D’Orazio, D. J., & Di Stefano, R. 2018, *MNRAS*, 474, 2975
 D’Orazio, D. J., & Duffell, P. C. 2021, *ApJL*, 914, L21
 D’Orazio, D. J., Haiman, Z., & Schiminovich, D. 2015, *Natur*, 525, 351
 Dotti, M., Colpi, M., Haardt, F., & Mayer, L. 2007, *MNRAS*, 379, 956
 Dressler, A. 1989, Observational Evidence for Supermassive Black Holes, Active Galactic Nuclei in International Astronomical Union. Symposium, 134, ed. D. E. Osterbrock & J. S. Miller (Dordrecht: Kluwer), 217
 Duffell, P. C., D’Orazio, D., Derdzinski, A., et al. 2020, *ApJ*, 901, 25
 Duncan, K., Conselice, C. J., Mundy, C., et al. 2019, *ApJ*, 876, 110
 Dvorkin, I., & Barausse, E. 2017, *MNRAS*, 470, 4547
 Ellis, J., & van Haasteren, R. 2017, jellis18/PTMCMCSampler: Official Release, v1.0.0, Zenodo, doi:10.5281/zenodo.1037579
 Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2020, 5ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, v3.0.0, Zenodo, doi:10.5281/zenodo.405981
 Enoki, M., Inoue, K. T., Nagashima, M., & Sugiyama, N. 2004, *ApJ*, 615, 19
 Enoki, M., & Nagashima, M. 2007, *PTPh*, 117, 241
 Eracleous, M., Boroson, T. A., Halpern, J. P., & Liu, J. 2012, *ApJS*, 201, 23
 Escala, A., Larson, R. B., Coppi, P. S., & Mardones, D. 2005, *ApJ*, 630, 152
 Event Horizon Telescope Collaboration et al. 2022, *ApJL*, 940, L12
 Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, *ApJL*, 875, L1
 Fabian, A. C. 2012, *ARA&A*, 50, 455
 Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2014, *ApJ*, 783, 134
 Finn, L. S., & Thorne, K. S. 2000, *PhRvD*, 62, 124021
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Gebhardt, K., Bender, R., Bower, G., et al. 2000, *ApJL*, 539, L13
 Gezari, S., Halpern, J. P., & Eracleous, M. 2007, *ApJS*, 169, 167
 Ghez, A. M., Klein, B. L., Morris, M., & Becklin, E. E. 1998, *ApJ*, 509, 678
 Goncharov, B., Shannon, R. M., Reardon, D. J., et al. 2021, *ApJL*, 917, L19
 Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, *MNRAS*, 453, 1562
 GRAVITY Collaboration, Abuter, R., Amorim, A., et al. 2018, *A&A*, 618, L10
 Gültekin, K., Richstone, D. O., Gebhardt, K., et al. 2009a, *ApJ*, 698, 198
 Gültekin, K., Richstone, D. O., Gebhardt, K., et al. 2009b, *ApJ*, 698, 198
 Haiman, Z., Kocsis, B., & Menou, K. 2009, *ApJ*, 700, 1952
 Häring, N., & Rix, H.-W. 2004, *ApJL*, 604, L89
 Hazboun, J. S., Simon, J., Taylor, S. R., et al. 2020, *ApJ*, 890, 108
 Hellings, R. W., & Downs, G. S. 1983, *ApJL*, 265, L39
 Hinton, S. R. 2016, *JOSS*, 1, 00045
 Hoffman, L., & Loeb, A. 2007, *MNRAS*, 377, 957
 Hogg, D. W. 1999, arXiv:astro-ph/9905116
 Holgado, A. M., Sesana, A., Sandrinelli, A., et al. 2018, *MNRAS*, 481, L74
 Holley-Bockelmann, K., & Khan, F. M. 2015, *ApJ*, 810, 139
 Holley-Bockelmann, K., & Sigurdsson, S. 2006, arXiv:astro-ph/0601520
 Hopkins, P. F., Hernquist, L., Cox, T. J., & Kereš, D. 2008, *ApJS*, 175, 356
 Huerta, E. A., McWilliams, S. T., Gair, J. R., & Taylor, S. R. 2015, *PhRvD*, 92, 063010
 Flanagan, É. É., & Hughes, S. A. 1998, *PhRvD*, 57, 4535
 Hunter, J. D. 2007, *CSE*, 9, 90
 Inayoshi, K., Ichikawa, K., & Haiman, Z. 2018, *ApJL*, 863, L36
 Jaffe, A. H., & Backer, D. C. 2003, *ApJ*, 583, 616
 Johnson, A. D., Vigeland, S. J., Siemens, X., & Taylor, S. R. 2022, *ApJ*, 932, 105
 Ju, W., Greene, J. E., Rafikov, R. R., Bickerton, S. J., & Badenes, C. 2013, *ApJ*, 777, 44
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
 Kelley, L., Charisi, M., Burke-Spolaor, S., et al. 2019a, *BAAS*, 51, 490
 Kelley, L. Z. 2021a, *MNRAS*, 500, 4065
 Kelley, L. Z. 2021b, *JOSS*, 6, 2784
 Kelley, L. Z., Blecha, L., & Hernquist, L. 2017a, *MNRAS*, 464, 3131
 Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2017b, *MNRAS*, 471, 4508
 Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2018, *MNRAS*, 477, 964

- Kelley, L. Z., Haiman, Z., Sesana, A., & Hernquist, L. 2019b, *MNRAS*, **485**, 1579
- Khan, F. M., Just, A., & Merritt, D. 2011, *ApJ*, **732**, 89
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, ed. F. Loizides & B. Schmidt (Netherlands: IOS Press), 87
- Kocsis, B., & Sesana, A. 2011, *MNRAS*, **411**, 1467
- Komossa, S. 2006, *MmSAI*, **77**, 733
- Kormendy, J. 1993, in *The Nearest Active Galaxies*, ed. J. Beckman, L. Colina, & H. Netzer, 197
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, **51**, 511
- Koss, M., Mushotzky, R., Treister, E., et al. 2012, *ApJL*, **746**, L22
- Kulier, A., Ostriker, J. P., Natarajan, P., Lackner, C. N., & Cen, R. 2015, *ApJ*, **799**, 178
- Lacey, C., & Cole, S. 1993, *MNRAS*, **262**, 627
- Lam, S. K., Pitrou, A., & Seibert, S. 2015, in *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC*, 1
- Lam, W. G., Taylor, S. R., & van Haasteren, R. 2023, in *Proc. of the Second Workshop on the LLVM Compiler Infrastructure in HPC, Numba: A LLVM-Based Python JIT Compiler* (New York: Association for Computing Machinery)
- Lang, P., Wuyts, S., Somerville, R. S., et al. 2014, *ApJ*, **788**, 11
- Liu, T., Gezari, S., Ayers, M., et al. 2019, *ApJ*, **884**, 36
- Magorrian, J., Tremaine, S., Richstone, D., et al. 1998, *AJ*, **115**, 2285
- Marconi, A., & Hunt, L. K. 2003, *ApJL*, **589**, L21
- McConnell, N. J., & Ma, C.-P. 2013, *ApJ*, **764**, 184
- McWilliams, S. T., Ostriker, J. P., & Pretorius, F. 2014, *ApJ*, **789**, 156
- Menou, K., Haiman, Z., & Narayanan, V. K. 2001, *ApJ*, **558**, 535
- Merritt, D., & Milosavljević, M. 2005, *LRR*, **8**, 8
- Middleton, H., Del Pozzo, W., Farr, W. M., Sesana, A., & Vecchio, A. 2016, *MNRAS*, **455**, L72
- Middleton, H., Sesana, A., Chen, S., et al. 2021, *MNRAS*, **502**, L99
- Milosavljević, M., & Merritt, D. 2001, *ApJ*, **563**, 34
- Mingarelli, C. M. F., Lazio, T. J. W., Sesana, A., et al. 2017, *NatAs*, **1**, 886
- Moody, M. S. L., Shi, J.-M., & Stone, J. M. 2019, *ApJ*, **875**, 66
- Mundy, C. J., Conzelmann, C. J., Duncan, K. J., et al. 2017, *MNRAS*, **470**, 3507
- Muñoz, D. J., Miranda, R., & Lai, D. 2019, *ApJ*, **871**, 84
- Nguyen, K., Bogdanović, T., Runnoe, J. C., et al. 2020, *ApJL*, **900**, L42
- Ostriker, J. P., & Hausman, M. A. 1977, *ApJL*, **217**, L125
- Peters, P. C. 1964, *PhRv*, **136**, 1224
- Peters, P. C., & Mathews, J. 1963, *PhRv*, **131**, 435
- Pflueger, B. J., Nguyen, K., Bogdanović, T., et al. 2018, *ApJ*, **861**, 59
- Phinney, E. S. 2001, arXiv:astro-ph/0108028
- Pol, N., Taylor, S. R., & Romano, J. D. 2022, *ApJ*, **940**, 173
- Pol, N. S., Taylor, S. R., Kelley, L. Z., et al. 2021, *ApJL*, **911**, L34
- Popović, L. Č. 2012, *NewAR*, **56**, 74
- Rajagopal, M., & Romani, R. W. 1995, *ApJ*, **446**, 543
- Rasskazov, A., & Merritt, D. 2017, *PhRvD*, **95**, 084032
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., & Hobbs, G. 2015, *MNRAS*, **447**, 2772
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., Hobbs, G., & Manchester, R. N. 2014, *MNRAS*, **442**, 56
- Richstone, D., Ajhar, E. A., Bender, R., et al. 1998, *Natur*, **385**, A14
- Rodríguez, C., Taylor, G. B., Zavala, R. T., et al. 2006, *ApJ*, **646**, 49
- Roebber, E., Holder, G., Holz, D. E., & Warren, M. 2016, *ApJ*, **819**, 163
- Rosado, P. A., & Sesana, A. 2014, *MNRAS*, **439**, 3986
- Rosado, P. A., Sesana, A., & Gair, J. 2015, *MNRAS*, **451**, 2417
- Runnoe, J. C., Eracleous, M., Mathes, G., et al. 2015, *ApJS*, **221**, 7
- Runnoe, J. C., Eracleous, M., Pennell, A., et al. 2017, *MNRAS*, **468**, 1683
- Ryu, T., Perma, R., Haiman, Z., Ostriker, J. P., & Stone, N. C. 2018, *MNRAS*, **473**, 3410
- Saglia, R. P., Opitsch, M., Erwin, P., et al. 2016, *ApJ*, **818**, 47
- Salcido, J., Bower, R. G., Theuns, T., et al. 2016, *MNRAS*, **463**, 870
- Sampson, L., Cornish, N. J., & McWilliams, S. T. 2015, *PhRvD*, **91**, 084055
- Saslaw, W. C., Valtonen, M. J., & Aarseth, S. J. 1974, *ApJ*, **190**, 253
- Sato-Polito, G., & Kamionkowski, M. 2023, arXiv:2305.05690
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521
- Schechter, P. 1976, *ApJ*, **203**, 297
- Sesana, A. 2010, *ApJ*, **719**, 851
- Sesana, A. 2013a, *CQGra*, **30**, 224014
- Sesana, A. 2013b, *MNRAS*, **433**, L1
- Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, *ApJ*, **611**, 623
- Sesana, A., Haiman, Z., Kocsis, B., & Kelley, L. Z. 2018, *ApJ*, **856**, 42
- Sesana, A., & Khan, F. M. 2015, *MNRAS*, **454**, L66
- Sesana, A., Shankar, F., Bernardi, M., & Sheth, R. K. 2016, *MNRAS*, **463**, L6
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *MNRAS*, **390**, 192
- Sesana, A., Vecchio, A., & Volonteri, M. 2009, *MNRAS*, **394**, 2255
- Shannon, R. M., Ravi, V., Coles, W. A., et al. 2013, *Sci*, **342**, 334
- Shannon, R. M., Ravi, V., Lentati, L. T., et al. 2015, *Sci*, **349**, 1522
- Shen, Y., Liu, X., Loeb, A., & Tremaine, S. 2013, *ApJ*, **775**, 49
- Simon, J. 2023, *ApJL*, **949**, L24
- Simon, J., & Burke-Spolaor, S. 2016, *ApJ*, **826**, 11
- Simon, J., Polin, A., Lommen, A., et al. 2014, *ApJ*, **784**, 60
- Siwek, M., Weinberger, R., & Hernquist, L. 2023, *MNRAS*, **522**, 2707
- Siwek, M. S., Kelley, L. Z., & Hernquist, L. 2020, *MNRAS*, **498**, 537
- Snyder, G. F., Lotz, J. M., Rodríguez-Gomez, V., et al. 2017, *MNRAS*, **468**, 207
- Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., & Hernquist, L. 2008, *MNRAS*, **391**, 481
- 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
- Taylor, S. R., & Gerosa, D. 2018, *PhRvD*, **98**, 083017
- Taylor, S. R., Simon, J., & Sampson, L. 2017, *PhRvL*, **118**, 181102
- Tomczak, A. R., Quadri, R. F., Tran, K.-V. H., et al. 2014, *ApJ*, **783**, 85
- Tremaine, S., Gebhardt, K., Bender, R., et al. 2002, *ApJ*, **574**, 740
- Tsalmantza, P., Decarli, R., Dotti, M., & Hogg, D. W. 2011, *ApJ*, **738**, 20
- Vallisneri, M., Taylor, S. R., Simon, J., et al. 2020, *ApJ*, **893**, 112
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, *MNRAS*, **461**, 3145
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014, *MNRAS*, **444**, 1518
- Volonteri, M., Haardt, F., & Madau, P. 2003, *ApJ*, **582**, 559
- Volonteri, M., Pfister, H., Beckmann, R. S., et al. 2020, *MNRAS*, **498**, 2219
- Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 3291
- White, S. D. M. 1980, *MNRAS*, **191**, 1P
- Wyithe, J. S. B., & Loeb, A. 2003, *ApJ*, **590**, 691
- Xin, C., Charisi, M., Haiman, Z., et al. 2020, *MNRAS*, **496**, 1683
- Yu, Q. 2002, *MNRAS*, **331**, 935
- Zhu, X.-J., Cui, W., & Thrane, E. 2019, *MNRAS*, **482**, 2588
- Zrake, J., Tiede, C., MacFadyen, A., & Haiman, Z. 2021, *ApJL*, **909**, L13