

Kalman tracking and parameter estimation of continuous gravitational waves with a pulsar timing array

Tom Kimpson,^{1,2,3}★ Andrew Melatos,^{1,2} Joseph O’Leary,^{1,2} Julian B. Carlin^{1,2} Robin J. Evans,⁴ William Moran,⁴ Tong Cheunchitra^{1,2} Wenhao Dong^{1,2} Liam Dunn^{1,2} Julian Greentree,³ Nicholas J. O’Neill,^{1,2} Sofia Suvorova,⁴ Kok Hong Thong^{1,2} and Andrés F. Vargas^{1,2}

¹*School of Physics, University of Melbourne, Parkville, VIC 3010, Australia*

²*ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav), University of Melbourne, Parkville, VIC 3010, Australia*

³*Department of Space and Climate Physics, University College London, Holmbury St. Mary RH5 6NT, UK*

⁴*Department of Electrical and Electronic Engineering, University of Melbourne, Parkville, VIC 3010, Australia*

Accepted 2024 September 19. Received 2024 September 15; in original form 2023 November 7

ABSTRACT

Continuous nanohertz gravitational waves from individual supermassive black hole binaries may be detectable with pulsar timing arrays. A novel search strategy is developed, wherein intrinsic achromatic spin wandering is tracked simultaneously with the modulation induced by a single gravitational wave source in the pulse times of arrival. A two-step inference procedure is applied within a state-space framework, such that the modulation is tracked with a Kalman filter, which then provides a likelihood for nested sampling. The procedure estimates the static parameters in the problem, such as the sky position of the source, without fitting for ensemble-averaged statistics such as the power spectral density of the timing noise, and therefore complements traditional parameter estimation methods. It also returns the Bayes factor relating a model with a single gravitational wave source to one without, complementing traditional detection methods. It is shown via astrophysically representative software injections in Gaussian measurement noise that the procedure distinguishes a gravitational wave from pure noise down to a characteristic wave strain of $h_0 \approx 2 \times 10^{-15}$. Full posterior distributions of model parameters are recovered and tested for accuracy. There is a bias of ≈ 0.3 rad in the marginalized one-dimensional posterior for the orbital inclination ι , introduced by dropping the so-called pulsar terms. Smaller biases $\lesssim 10$ per cent are also observed in other static parameters.

Key words: gravitational waves – pulsars: general.

1 INTRODUCTION

The inspiral of supermassive black hole binaries (SMBHBs; Rajagopal & Romani 1995; Jaffe & Backer 2003; Wyithe & Loeb 2003; Sesana 2013; McWilliams, Ostriker & Pretorius 2014; Ravi et al. 2015; Burke-Spolaor et al. 2019; Sykes et al. 2022) is predicted to emit nanohertz (nHz) gravitational waves (GWs). Other GW sources in this low-frequency regime include cosmic strings (e.g. Sanidas, Battye & Stappers 2012) and cosmological phase transitions (e.g. Xue et al. 2021). The detection of nHz GWs has inspired the development of new observational methods, since it is impractical to engineer terrestrial interferometric detectors with sufficiently long baselines. The foremost method is timing an ensemble of pulsars, that is, a pulsar timing array (PTA; Tiburzi 2018; Verbiest, Osłowski & Burke-Spolaor 2021). A nHz GW influences the trajectory and frequency of individual radio pulses, leaving a characteristic impression on the pulse times of arrival (TOAs) measured at the Earth. By measuring TOAs from multiple pulsars simultaneously one can effectively construct a detector with a baseline on the scale of parsecs. Multiple

PTA detectors have been built over the last few decades, including the North American Nanohertz Observatory for Gravitational Waves (NANOGrav, Agazie et al. 2023b), the Parkes Pulsar Timing Array (PPTA, Zic et al. 2023), and the European Pulsar Timing Array (EPTA, Antoniadis et al. 2023b). These individual efforts have joined in international collaboration, under the umbrella of the International Pulsar Timing Array (IPTA, Perera et al. 2019), along with a number of newer PTAs such as the Indian Pulsar Timing Array Project (InPTA, Tarafdar et al. 2022), MeerTime (Bailes et al. 2020; Spiewak et al. 2022), and the Chinese PTA (CPTA, Hobbs et al. 2019).

The incoherent superposition of multiple SMBHB sources leads to a stochastic GW background at nHz frequencies (Allen 1997; Sesana, Vecchio & Colacino 2008; Christensen 2019; Renzini et al. 2022). Previous efforts have mainly focused on detecting the stochastic background by measuring the cross-correlation between the timing residuals from pairs of pulsars as a function of the angular separation between the pulsars – the Hellings–Downs curve (Hellings & Downs 1983). After multiple non-detections (Lentati et al. 2015; Arzoumanian et al. 2018b; Antoniadis et al. 2022) consistent evidence for the GW background was presented by NANOGrav (Agazie et al. 2023a), EPTA/InPTA (Antoniadis et al. 2023a), PPTA (Reardon et al. 2023), and the CPTA (Xu et al. 2023).

* E-mail: tom.kimpson@unimelb.edu.au

Individual SMBHBs that are sufficiently massive and nearby may be resolvable with PTAs, allowing the early stages of their evolution and coalescence to be investigated (Sesana & Vecchio 2010; Yardley et al. 2010; Babak & Sesana 2012; Ellis 2013; Zhu et al. 2015a, 2016). Indeed, the stochastic GW background itself may be dominated by a few individual binary sources (Ravi et al. 2012). Individual SMBHBs are continuous wave sources; they generate persistent, quasi-monochromatic modulations of a known form in pulsar timing residuals. Consequently, they are detected more efficiently by either a frequentist matched filter, for example, the \mathcal{F} -statistic (Lee et al. 2011; Ellis, Siemens & Creighton 2012; Zhu et al. 2014), or else Bayesian inference (Ellis & Cornish 2016; Arzoumanian et al. 2020a), rather than by cross-correlating pulsar pairs. However, PTA observational campaigns to detect individual sources have been unsuccessful so far (Jenet et al. 2004; Zhu et al. 2014; Babak et al. 2016; Arzoumanian et al. 2023). Inconclusive evidence at low significance was presented recently by the EPTA for an individual source at 4–5 nHz (Antoniadis et al. 2023c).

Intrinsic pulsar timing noise – that is, random, unmodelled, red-spectrum TOA fluctuations due to irregularities in the rotation of the star – has been identified as a key factor limiting the sensitivity of PTAs to GW signals (Shannon & Cordes 2010; Lasky et al. 2015; Caballero et al. 2016; Goncharov et al. 2021). This timing noise has multiple theorized causes including free precession (Stairs, Lyne & Shemar 2000; Kerr et al. 2015), microglitches (D’Alessandro et al. 1995; Melatos, Peralta & Wyithe 2008; Espinoza et al. 2021), asteroid encounters (Brook et al. 2013; Shannon et al. 2013), glitch recovery (Johnston & Galloway 1999; Hobbs, Lyne & Kramer 2010), fluctuations in internal and external stochastic torques (Cordes & Greenstein 1981; Urama, Link & Weisberg 2006; Meyers, Melatos & O’Neill 2021a; Meyers et al. 2021b; Antonelli, Basu & Haskell 2023), variations in the coupling between the stellar crust and core (Jones 1990; Meyers et al. 2021b; Melatos et al. 2023), magnetospheric state switching (Kramer et al. 2006; Lyne et al. 2010; Stairs et al. 2019), and superfluid turbulence (Greenstein 1970; Peralta et al. 2006; Melatos & Link 2014). In order to mitigate the impact of timing noise, PTAs are typically composed of millisecond pulsars (MSPs), which are relatively stable rotators. However, timing noise in MSPs may be a latent phenomenon that will increasingly assert itself as longer stretches of more sensitive data are analysed in the quest to detect nHz GWs (Shannon & Cordes 2010). In modern Bayesian PTA searches, the power spectral density of the red intrinsic timing noise is modelled (usually as a broken or unbroken power law) and estimated, in an effort to distinguish it from the noise induced by a stochastic GW background (whose spectrum is also red). In addition to the red timing noise there are secondary, white noise sources to consider such as phase jitter noise and radiometer noise (Cordes & Shannon 2010; Lam et al. 2019; Parthasarathy et al. 2021).

In this work, we present an alternative and complementary approach to PTA data analysis for individual, quasi-monochromatic, SMBHB sources which self-consistently tracks the intrinsic timing noise in PTA pulsars and disentangles it from GW-induced TOA modulations. The new approach differs from existing approaches in one key respect: it infers the GW parameters conditional on the unique, time-ordered realization of the noisy TOAs observed, instead of fitting for the ensemble-averaged statistics of the TOA noise process, for example, the amplitude and exponent of its power spectral density. Stated another way, existing approaches seek to detect a GW signal by marginalizing over the ensemble of possible noise realizations summarized by the power spectral density, whereas the new approach delivers the most likely set of GW parameters consistent with the actual, observed noise realization. The new and

existing approaches are therefore complementary. In particular, we formulate PTA analysis as a state-space problem and demonstrate how to optimally estimate the state-space evolution using a Kalman filter, a tried-and-tested tool (Kalman 1960; Meyers et al. 2021b; Melatos et al. 2023). We combine the Kalman tracking of the pulsars’ intrinsic rotational states with a Bayesian nested sampler (Skilling 2006; Ashton et al. 2022) to estimate the static GW parameters and calculate the marginal likelihood (i.e. the model evidence) for model selection. The adaptive bias of the Kalman filter tracks timing noise more effectively than alternative techniques such as least-squares estimators.

The paper is organized as follows. In Section 2, we present the state-space model for the rotational states of an array of pulsars falling freely in the curved space–time of a single-source GW. In Section 3, we develop a Kalman filter to track the state evolution and deploy the Kalman filter in conjunction with nested sampling to estimate the GW and other system parameters, along with the model evidence. In Section 4, we describe how we create synthetic validation data to test the method. In Section 5, we test the method on synthetic data for a single representative GW source. In Section 6, we extend the tests to cover an astrophysically relevant domain of SMBHB source parameters. In Section 7, we quantify the bias in the parameter estimates. In Section 8, we review the computational cost of the method. Conclusions are drawn in Section 9. The data are formulated as pulse frequency time-series with Gaussian measurement noise as a proof of principle and to maintain consistency with previous work (Meyers et al. 2021a, b). It will be necessary to modify the method to accept pulse TOAs instead of a pulse frequency time-series when analysing real, astronomical data, a subtle generalization which is deferred to future work. Throughout the paper we adopt natural units, with $c = G = \hbar = 1$, and metric signature $(-, +, +, +)$.

2 STATE-SPACE FORMULATION

We formulate the PTA analysis as a state-space problem, in which the intrinsic rotational state of each pulsar evolves according to a stochastic differential equation and is related to the observed pulse sequence via a measurement equation. In this work, we take the intrinsic state variable to be the n th pulsar’s spin frequency $f_p^{(n)}(t)$, as measured in the local, freely falling rest frame of the pulsar’s centre of mass. A phenomenological model for the evolution of $f_p^{(n)}(t)$ is presented in Section 2.1. We take the measurement variable to be the radio pulse frequency measured by an observer at Earth, $f_m^{(n)}(t)$. The measurement equation relating $f_m^{(n)}(t)$ to $f_p^{(n)}(t)$ is presented in Section 2.2. The superscript $1 \leq n \leq N$ indexes the n th pulsar in the array. The subtle problem of generalizing the measurement variable to pulse TOAs is postponed to future work, as noted in Introduction.

2.1 Spin evolution

A predictive, first-principles theory of timing noise does not exist at present; there are several plausible physical mechanisms, referenced in Introduction. We therefore rely on an idealized phenomenological model to capture the main qualitative features of a typical PTA pulsar’s observed spin evolution, that is, random, mean-reverting, small-amplitude excursions around a smooth, secular trend. In the model, $f_p^{(n)}(t)$ evolves according to the sum of a deterministic and a stochastic torques. The deterministic torque is attributed to magnetic dipole braking, with braking index $n_{\text{em}} = 3$ for the sake of definiteness (Goldreich & Julian 1969). Most PTAs involve MSPs, for which the quadratic correction due to n_{em} in $f_p^{(n)}(t)$ is negligible over the observation time $T_{\text{obs}} \sim 10$ yr, and the deterministic evolution $f_{\text{em}}^{(n)}(t)$

can be approximated accurately by

$$f_{\text{em}}^{(n)}(t) = f_{\text{em}}^{(n)}(t_1) + \dot{f}_{\text{em}}^{(n)}(t_1)t, \quad (1)$$

where an overdot denotes a derivative with respect to t , and t_1 labels the time of the first TOA. The stochastic torque is assumed to be a zero-mean, white noise process. Specifically, the frequency evolves according to an Ornstein–Uhlenbeck process, described by a Langevin equation with a time-dependent drift term (Vargas & Melatos 2023),

$$\frac{df_p^{(n)}}{dt} = -\gamma^{(n)}[f_p^{(n)} - f_{\text{em}}^{(n)}(t)] + \dot{f}_{\text{em}}^{(n)}(t) + \xi^{(n)}(t). \quad (2)$$

In equation (2), $f_{\text{em}}^{(n)}$ is the solution of the electromagnetic spin-down equation given by equation (1), $\dot{f}_{\text{em}}^{(n)}$ is the spin derivative, $\gamma^{(n)}$ is a damping constant whose reciprocal specifies the mean-reversion time-scale, and $\xi^{(n)}(t)$ is a white noise stochastic process which satisfies

$$\langle \xi^{(n)}(t) \rangle = 0, \quad (3)$$

$$\langle \xi^{(n)}(t) \xi^{(n')}(t') \rangle = [\sigma^{(n)}]^2 \delta_{n,n'} \delta(t - t'). \quad (4)$$

In equation (4) $[\sigma^{(n)}]^2$ is the variance of $\xi^{(n)}$ and parametrizes the amplitude of the noise. Combined with the mean reversion it gives characteristic root mean square fluctuations $\approx \sigma^{(n)}/[\gamma^{(n)}]^{1/2}$ in $f_p^{(n)}(t)$ (Gardiner 2009). It is important to note that white noise fluctuations in $\xi(t)$ translate into red noise fluctuations in the rotational phase $\phi(t) = \int_{t_1}^t dt' f_p(t')$ after being filtered by the terms involving d/dt and $\gamma^{(n)}$ in equation (2), consistent with the observed power spectral density of typical MSPs in the nHz band relevant to PTA experiments.

Equations (1)–(4) represent a phenomenological model, which aims to reproduce qualitatively the typical timing behaviour observed in PTAs, viz. a mean-reverting random walk about a secular spin-down trend (Agazie et al. 2023b; Antoniadis et al. 2023b; Zic et al. 2023). Equations (1)–(4) are not derived from first principles by applying a microphysical theory. As a first pass, they also exclude certain phenomenological elements, which are likely to be present in reality, for example, the classic, two-component, crust-superfluid structure inferred from post-glitch recoveries (Baym et al. 1969; van Eysden & Melatos 2010; Gügercinoğlu & Alpar 2017; Meyers et al. 2021a, b). An approach akin to equations (1)–(4) has been followed successfully in other timing analyses in the context of anomalous braking indices (Vargas & Melatos 2023) and hidden Markov model glitch searches (Melatos et al. 2020; Lower et al. 2021; Dunn et al. 2022, 2023). However, equations (1)–(4) involve significant idealizations, which must be recognized at the outset (Meyers et al. 2021a, b; Vargas & Melatos 2023). First, the white noise driver $\xi(t)$ in equation (2) is not differentiable, which makes the formal interpretation of $d^2 f_p/dt^2$ ambiguous, even though $d^2 f_p/dt^2$ is not used in the PTA analysis proposed in this paper. Second, the white spectrum assumed for $\xi(t)$ may or may not be suitable for MSPs in PTAs. It is challenging observationally to infer the spectrum of $\xi(t)$ from the observed spectrum of the phase residuals, because the inference is conditional on the (unknown) dynamical model governing $d f_p/dt$. For small-amplitude fluctuations sampled relatively often, as in MSPs in PTAs, it is likely that $\xi(t)$ is white to a good approximation over the inter-TOA intervals and generates red phase residuals as observed, but caution is warranted nevertheless. Third, the Brownian increment $d\mathcal{B}(t) = \xi(t)dt$ does not include non-Gaussian excursions such as Lévy flights (Sornette 2004), which have not been ruled out by pulsar timing experiments to date. The above three idealizations are supplemented by other, physical

approximations noted above, for example, neglecting n_{em} in equation (1) and differential rotation between the crust and superfluid in equation (2).

2.2 Modulation of pulsar frequency by a GW

In the presence of a GW, the pulse frequency measured by an observer in the local rest frame of the neutron star’s centre of mass is different from that measured by an observer on Earth. Specifically, the pulse frequency at the Earth is modulated harmonically at the GW frequency. We derive the non-linear measurement equation relating $f_{\text{m}}(t)$ to $f_p(t)$ in this section. The measurement equation is a key input into the Kalman filter in Section 3.1

2.2.1 Plane GW perturbation

We consider a gravitational plane wave from a single, distant SMBHB. The GW perturbs the background Minkowski metric $\eta_{\mu\nu}$ as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (5)$$

where the metric perturbation $h_{\mu\nu}$ has zero temporal components $h_{0\nu} = h_{\mu 0} = 0$. For elliptically polarized GWs emitted by a SMBHB, the spatial metric components are (Maggiore 2018)

$$h_{ij}(t, \mathbf{x}) = H_{ij}^{(+)} e^{i[\Omega(\mathbf{x} \cdot \mathbf{n} - t) + \Phi_0]} + H_{ij}^{(\times)} e^{i[\Omega(\mathbf{x} \cdot \mathbf{n} - t) + \Phi_0 + \pi/2]}, \quad (6)$$

written in terms of nearly Lorentz spatial coordinates and global coordinate x time t (Schutz 2022). The GW propagates in the \mathbf{n} -direction (where \mathbf{n} is a unit vector), has a constant (see justification below) angular frequency Ω , phase offset Φ_0 , and two orthogonal polarizations with amplitude tensors $H_{ij}^{(+,\times)}$. Throughout this paper, we work with pulsar TOAs defined relative to the Solar system barycentre (SSB). We are free to choose coordinates such that Φ_0 is the GW phase at $t = 0$ at the SSB. The amplitude tensors are given by

$$H_{ij}^{(+)} = h_+ e_{ij}^+, \quad (7)$$

$$H_{ij}^{(\times)} = h_\times e_{ij}^\times, \quad (8)$$

where h_+ and h_\times are the respective polarization amplitudes. The plus and cross-polarization tensors e_{ij}^+ and e_{ij}^\times are uniquely defined by the principal axes of the wave, viz. the unit three-vectors \mathbf{k} and \mathbf{l} , according to

$$e_{ij}^+(\mathbf{n}) = k_i k_j - l_i l_j, \quad (9)$$

$$e_{ij}^\times(\mathbf{n}) = k_i l_j + l_i k_j. \quad (10)$$

The principal axes are in turn specified by the location of the GW source on the sky (colatitude θ , azimuth ϕ) and the polarization angle ψ according to

$$\begin{aligned} \mathbf{k} = & (\sin \phi \cos \psi - \sin \psi \cos \phi \cos \theta) \hat{\mathbf{x}} \\ & - (\cos \phi \cos \psi + \sin \psi \sin \phi \cos \theta) \hat{\mathbf{y}} \\ & + (\sin \psi \sin \theta) \hat{\mathbf{z}}, \end{aligned} \quad (11)$$

$$\begin{aligned} \mathbf{l} = & (-\sin \phi \sin \psi - \cos \psi \cos \phi \cos \theta) \hat{\mathbf{x}} \\ & + (\cos \phi \sin \psi - \cos \psi \sin \phi \cos \theta) \hat{\mathbf{y}} \\ & + (\cos \psi \sin \theta) \hat{\mathbf{z}}, \end{aligned} \quad (12)$$

where, for example, $\hat{\mathbf{x}}$ is a unit vector in the direction of the x -axis. The direction of GW propagation is related to the principal axes by

$$\mathbf{n} = \mathbf{k} \times \mathbf{l}. \quad (13)$$

In this paper, the source is approximated as monochromatic. In reality, the GW frequency $f_{\text{gw}} = \Omega/2\pi$ increases during the inspiral by an amount (e.g. Sesana & Vecchio 2010)

$$\Delta f_{\text{gw}} = 0.05 \text{ nHz} \left(\frac{M_c}{10^{8.5} M_\odot} \right)^{5/3} \left[\frac{f_{\text{gw}}(t = t_1)}{50 \text{ nHz}} \right]^{11/3} \left(\frac{T_{\text{obs}}}{10 \text{ yr}} \right) \quad (14)$$

where M_c is the chirp mass of the SMBHB, $f_{\text{gw}}(t = t_1)$ is the GW frequency at the time of the first observation, and T_{obs} is the length of time over which Δf_{gw} is measured, which for PTAs is ~ 10 yr. A source can be considered monochromatic, if Δf_{gw} is less than the PTA frequency resolution $1/T_{\text{obs}}$. Equation (14) implies $(M_c/10^{8.5} M_\odot)^{5/3} [f_{\text{gw}}(t = t_1)/50 \text{ nHz}]^{11/3} (T_{\text{obs}}/10 \text{ yr})^2 \gtrsim 20$ for $\Delta f_{\text{gw}} \gtrsim 1/T_{\text{obs}}$. The majority of SMBHBs detectable with PTAs are expected to satisfy $\Delta f_{\text{gw}} < 1/T_{\text{obs}}$; for a PTA composed of pulsars with a mean distance of 1.5 kpc, 78 per cent of simulated SMBHBs satisfy this condition for the current IPTA, whilst for the second phase of the Square Kilometer Array this fraction drops to 52 per cent; see fig. 7 in Rosado, Sesana & Gair (2015). We are therefore justified in treating the GW source as monochromatic as a first pass in this introductory paper (Sesana et al. 2008; Sesana & Vecchio 2010; Ellis et al. 2012).

2.2.2 Measurement equation

In general radio pulses from a pulsar are transmitted as amplitude modulations of a radio-frequency carrier wave. They are described by the geometric object \mathbf{p} , which we identify as the momentum four-vector of the radio pulse. The presence of a GW induces a shift in the temporal component of the associated momentum one form, while the photon travels from the emitter to the observer, that is, $\Delta p_t = p_t|_{\text{observer}} - p_t|_{\text{emitter}}$. One obtains (e.g. Maggiore 2018)

$$\Delta p_t = \frac{\pi f_p h_{ij}(t; \mathbf{x} = 0) q^i q^j}{1 + \mathbf{n} \cdot \mathbf{q}} [1 - e^{i\Omega(1 + \mathbf{n} \cdot \mathbf{q})d}] , \quad (15)$$

where f_p is the pulse frequency measured in the momentarily comoving reference frame of an observer at rest in the nearly Lorentz coordinates (t, \mathbf{x}) . In equation (15), \mathbf{q} is the unit vector connecting the observer and the pulsar and d is the distance to the pulsar. We take the pulsar location to be constant, that is, neither \mathbf{q} nor d are functions of time. In practice, the pulsar locations vary with respect to the Earth but are constant with respect to the SSB. The barycentring correction is typically applied when generating TOAs, for example, with TEMPO2 (Hobbs, Edwards & Manchester 2006) and related timing software, and is inherited by the frequency time-series. Some pulsars, including some PTA pulsars, do have non-negligible proper motions of order 10^2 km s^{-1} after the barycentring corrections have been applied (e.g. Jankowski et al. 2018), but we do not consider this effect in this paper.

Generally, the measured frequency of a photon recorded by an observer who is travelling with four-velocity \mathbf{u} is given by the coordinate-independent expression $p_\alpha u^\alpha$. After barycentring, one has $u^\alpha = (1, 0, 0, 0)$ for both the emitter and the observer to leading order in the respective momentarily comoving reference frames. Hence, equation (15) can be written for the n th pulsar as

$$f_{\text{m}}^{(n)}(t) = f_p^{(n)} [t - d^{(n)}] g^{(n)}(t) + \varepsilon^{(n)}(t) , \quad (16)$$

where $d^{(n)}$ labels the distance to the n th pulsar, and $\varepsilon^{(n)}$ is a Gaussian measurement noise which satisfies

$$\langle \varepsilon^{(n)}(t) \rangle = 0 , \quad (17)$$

$$\langle \varepsilon^{(n)}(t) \varepsilon^{(n')}(t') \rangle = \sigma_{\text{m}}^2 \delta_{n,n'} \delta(t - t') , \quad (18)$$

where σ_{m} is the covariance of the measurement noise at the telescope and is shared between all pulsars by assumption. The measurement function $g^{(n)}(t)$ is related to a redshift $z^{(n)}(t)$ through

$$g^{(n)}(t) = 1 - z^{(n)}(t) , \quad (19)$$

with

$$z(t) = \frac{[q^{(n)}]^i [q^{(n)}]^j}{2[1 + \mathbf{n} \cdot \mathbf{q}^{(n)}]} \times \left[h_{ij}(t, \mathbf{x} = 0) - h_{ij}(t, \mathbf{x} = 0) e^{i\Omega(1 + \mathbf{n} \cdot \mathbf{q}^{(n)})d^{(n)}} \right] , \quad (20)$$

where $[q^{(n)}]^i$ labels the i th coordinate component of the n th pulsar's position vector $\mathbf{q}^{(n)}$. It is also instructive to express equation (20) in a trigonometric form (cf. e.g. Sesana & Vecchio 2010; Perrodin & Sesana 2018; Agazie et al. 2023c) as

$$z(t) = \frac{[q^{(n)}]^i [q^{(n)}]^j}{2[1 + \mathbf{n} \cdot \mathbf{q}^{(n)}]} \times \left\{ \begin{aligned} & \left[H_{ij}^{(+)} \cos \Phi(t) + H_{ij}^{(\times)} \sin \Phi(t) \right] \\ & - \left[H_{ij}^{(+)} \cos \Phi^{(n)}(t) + H_{ij}^{(\times)} \sin \Phi^{(n)}(t) \right] \end{aligned} \right\} , \quad (21)$$

where we define the phases

$$\Phi(t) = -\Omega t + \Phi_0 , \quad (22)$$

$$\Phi^{(n)}(t) = \Phi(t) + \Omega [1 + \mathbf{n} \cdot \mathbf{q}^{(n)}] d^{(n)} . \quad (23)$$

Equations (16)–(23) define a non-linear measurement equation that relates the intrinsic pulsar spin frequency to the pulse frequency measured by an observer on Earth.

3 SIGNAL TRACKING, PARAMETER ESTIMATION, AND MODEL SELECTION

The set of static parameters $\boldsymbol{\theta}$ of the model outlined in Section 2 can be separated into parameters controlling the intrinsic frequency evolution of the pulsars in the array and the GW source, that is,

$$\boldsymbol{\theta} = \boldsymbol{\theta}_{\text{psr}} \cup \boldsymbol{\theta}_{\text{gw}} , \quad (24)$$

with

$$\boldsymbol{\theta}_{\text{psr}} = \{\gamma^{(n)}, \sigma^{(n)}, f_{\text{em}}^{(n)}(t_1), f_{\text{em}}^{(n)}(t_1), d^{(n)}\}_{1 \leq n \leq N} , \quad (25)$$

and

$$\boldsymbol{\theta}_{\text{gw}} = \{h_0, \iota, \delta, \alpha, \psi, \Omega, \Phi_0\} , \quad (26)$$

where δ , α , and ι are the declination, right ascension, and inclination of the GW source respectively.¹ In equation (26), we reparametrize the two GW polarization amplitudes, h_+ and h_\times , in terms of ι and the characteristic wave strain h_0 through

$$h_+ = h_0(1 + \cos^2 \iota) , \quad (27)$$

$$h_\times = -2h_0 \cos \iota . \quad (28)$$

We use the parametrization in terms of h_0 and ι throughout the remainder of this work. A PTA containing N pulsars comprises $7 + 5N$ parameters to estimate. Typically the pulsar parameters are constrained better a priori by electromagnetic observations than the GW parameters. For example, estimates of pulsar distances are

¹ ι is the angle between the unit normal to the SMBHB orbital plane, \mathbf{L} , and the observer's line of sight, that is, $\cos \iota = \mathbf{n} \cdot \mathbf{L}$.

accurate to ~ 10 per cent (Cordes & Lazio 2002; Verbiest et al. 2012; Desvignes et al. 2016; Yao, Manchester & Wang 2017), but we have no prior information about δ and α .

In this section, we present a new method to infer θ and calculate the marginal likelihood (i.e. the model evidence). In Section 3.1, we outline how noisy measurements of the pulsar frequency, $f_m^{(n)}(t)$, can be used to estimate the hidden state sequence, $f_p^{(n)}(t)$, using a Kalman filter. In Section 3.2, we demonstrate how to deploy the Kalman filter in conjunction with a nested sampling technique to perform Bayesian inference. Model selection and the specification of the null model are described in Section 3.3. A complete summary of the workflow is presented in Section 3.4. The method complements traditional PTA analyses because (i) it harnesses the favourable asymptotic properties of the adaptive gain of the Kalman filter to track timing noise more nimbly than alternatives like least-squares estimators; and (ii) it infers θ conditional on the specific, time-ordered, random realization of the noise corresponding to the observed sequence of TOAs rather than on ensemble-averaged quantities like the power spectral density. Points (i) and (ii) are discussed further in Section 3.5.

3.1 Kalman filter and likelihood

The Kalman filter (1960) is a Gauss–Markov model used to algorithmically recover a temporal sequence of stochastically evolving system state variables, $X(t)$, which are not observed directly, given a temporal sequence of noisy measurements, $Y(t)$. It finds common use in engineering applications and has been applied successfully in neutron star astrophysics (e.g. Meyers et al. 2021a, b; Melatos et al. 2023). In this work we use the linear Kalman filter, which assumes a linear relation between dX/dt and $X(t)$ (dynamics) and between $Y(t)$ and $X(t)$ (measurement), with $X(t) = \{f_p^{(n)}(t)\}$ and $Y(t) = \{f_m^{(n)}(t)\}$. Extension to non-linear problems is straightforward using either an extended Kalman filter (Zarchan & Musoff 2000), unscented Kalman filter (Wan & Van Der Merwe 2000), or particle filter (Simon 2006). Equations (16) and (21) are non-linear in the static parameters (e.g. $d^{(n)}$), even though they are linear in $X(t)$ and $Y(t)$. Hence, inferring the static parameters is a non-linear exercise, to be tackled separately after the linear Kalman filter operates on the data for a fixed set of static parameters. Inference of the static parameters in equations (16) and (21) by nested sampling is discussed in Section 3.2.

Implementation instructions for the linear Kalman filter for the PTA state-space model in Section 2, including the full set of Kalman recursion relations, are presented in Appendix A. At each discrete time-step indexed by $1 \leq i \leq M$, the Kalman filter returns an estimate of the state variables, $\hat{X}_i = \hat{X}(t_i)$, and the covariance of those estimates, $P_i = \langle \hat{X}_i \hat{X}_i^T \rangle$, where the superscript T denotes the matrix transpose. The filter tracks the error in its predictions of X_i by converting \hat{X}_i into predicted measurements \hat{Y}_i via equations (16) and (21) and comparing with the actual, noisy measurements \hat{Y}_i . This defines a residual $\epsilon_i = Y_i - \hat{Y}_i$, which is sometimes termed the innovation. The Kalman filter also calculates the uncertainty in ϵ_i via the innovation covariance $S_i = \langle \epsilon_i \epsilon_i^T \rangle$. The innovation and the innovation covariance are then used to correct the state estimate \hat{X}_i according to equation (A8). For a fixed set of static parameters, the Kalman filter returns an estimate of the state sequence $\hat{X}_1, \dots, \hat{X}_M$ which minimizes the mean square error. We explain how to use this intermediate output to infer the optimal values of the static parameters θ in Section 3.2.

The Gaussian log-likelihood of obtaining Y_i given \hat{X}_i can then be calculated at each time-step from the Kalman filter output according

to

$$\log \mathcal{L}_i = -\frac{1}{2} (N \log 2\pi + \log |S_i| + \epsilon_i^T S_i^{-1} \epsilon_i) . \quad (29)$$

The total log-likelihood for the entire sequence is

$$\log \mathcal{L} = \sum_{i=1}^M \log \mathcal{L}_i . \quad (30)$$

Given Y_1, \dots, Y_M , \mathcal{L} is a function of the estimates $\hat{\theta}$ of the static parameters passed to the Kalman filter, that is, $\mathcal{L} = \mathcal{L}(Y|\hat{\theta})$. Similarly the estimates of the state and measurement variables, \hat{X} and \hat{Y} , are functions of $\hat{\theta}$. If $\hat{\theta}$ is close to the true underlying parameters θ , then the errors in \hat{X} and \hat{Y} are minimized and \mathcal{L} is maximized. This is illustrated with synthetic data in Fig. 1. In the left-hand column, a time-series of $f_m^{(1)}(t)$ including Gaussian noise (middle panel, red curve) is generated from equations (1)–(4), (16), and (21) for a single pulsar and fed into the Kalman filter along with the true static parameters $\hat{\theta} = \theta$. The Kalman filter recovers the evolution of $f_p^{(1)}(t)$ with high fidelity; the estimate of $f_p^{(1)}(t)$ (left top panel, blue curve) overlaps almost perfectly with the true $f_p^{(1)}(t)$ (left top panel, green curve). The predicted state $\hat{f}_p^{(1)}(t)$ is converted into a predicted measurement $\hat{f}_m^{(1)}(t)$ (middle panel, magenta curve), which again overlaps almost perfectly with the true measurement. The residuals $\epsilon(t)$ between the true and predicted measurements are small ($\lesssim 0.1$ per cent) and normally distributed (left bottom panel). By contrast, in the right-hand column, the exercise is repeated while passing incorrect static parameters ($\hat{\theta} \neq \theta$) to the Kalman filter, where Ω is displaced from its true value by 20 per cent. In this case, the Kalman filter fails to track $f_p^{(1)}(t)$ accurately, as the discrepancy between the blue and green curves in the top panel of the right-hand column indicates. It similarly fails to predict $f_m^{(1)}(t)$ accurately, as shown by the discrepancy between the red and magenta curves in the middle panel, and the residuals are no longer distributed normally (right bottom panel). In Section 3.2, we explain how to combine the Kalman filter with a nested sampler to iteratively guide $\hat{\theta}$ towards the true value of θ .

3.2 Nested sampling

We can use the likelihood returned by the Kalman filter, equation (30), in conjunction with likelihood-based inference methods to estimate the posterior distribution of θ by Bayes' rule,

$$p(\theta|Y) = \frac{\mathcal{L}(Y|\theta)\pi(\theta)}{\mathcal{Z}} , \quad (31)$$

where $\pi(\theta)$ is the prior distribution on θ and \mathcal{Z} is the marginalized likelihood, or evidence

$$\mathcal{Z} = \int d\theta \mathcal{L}(Y|\theta)\pi(\theta) . \quad (32)$$

We estimate the posterior distribution and the model evidence through nested sampling (Skilling 2006) in this paper. Nested sampling evaluates marginalized-likelihood integrals, of the form given by equation (32). It also approximates the posterior by returning samples from $p(\theta|Y)$. It does so by drawing a set of n_{live} live points from $\pi(\theta)$ and iteratively replacing the live point with the lowest likelihood with a new live point drawn from $\pi(\theta)$, where the new live point is required to have a higher likelihood than the discarded point. The primary advantage of nested sampling is its ability to compute \mathcal{Z} , on which model selection relies. Nested sampling is also computationally efficient and can handle multimodal

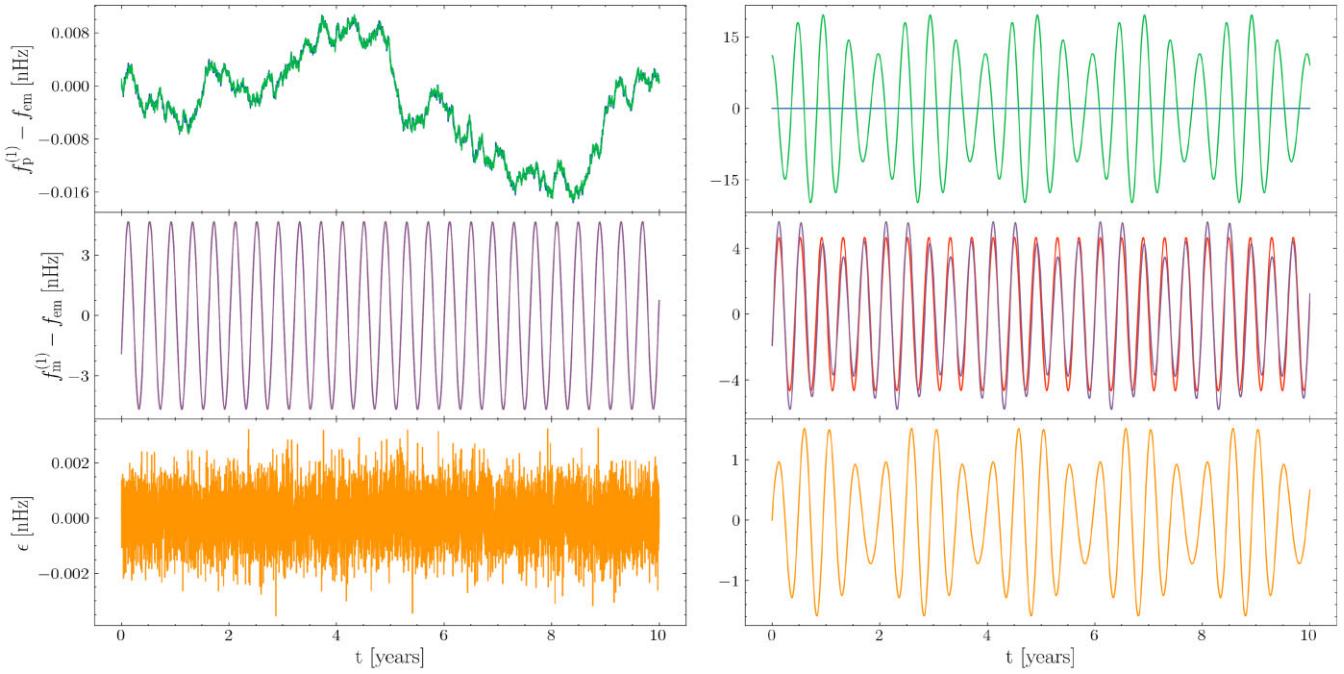


Figure 1. Sample output of the Kalman filter illustrating the accuracy of the reconstructed state sequence $f_p^{(1)}(t)$ when the static parameters are correct ($\hat{\theta} = \theta$, left-hand column) and incorrect ($\hat{\theta} \neq \theta$, right-hand column). The top panels show the true pulsar state $f_p^{(1)}(t) - f_{\text{em}}^{(1)}(t)$ (blue curve) and the state estimated by the Kalman filter $\hat{f}_p^{(1)}(t) - f_{\text{em}}^{(1)}(t)$ (green curve). We subtract $f_{\text{em}}^{(1)}(t)$ to better illustrate the stochastic wandering of the pulsar frequency. In the left-hand column, the blue/green solutions overlap almost perfectly; in the right-hand column, they do not. The middle panels show the true measured frequency $f_m^{(1)}(t) - f_{\text{em}}^{(1)}(t)$ (red curve) and the frequency estimated by the Kalman filter $\hat{f}_m^{(1)}(t) - f_{\text{em}}^{(1)}(t)$ (magenta curve). Again we subtract $f_{\text{em}}^{(1)}(t)$. In the left-hand column, the red/magenta solutions overlap almost perfectly; in the right-hand column, they do not. The bottom panels show the residual or innovation $\epsilon(t) = f_m^{(1)}(t) - \hat{f}_m^{(1)}(t)$. In the right-hand column, Ω is displaced from its true value by 20 per cent. Results are shown for a single pulsar.

problems (Ashton et al. 2022). For these reasons, it has enjoyed widespread adoption in the physical sciences, particularly within the cosmological community (Mukherjee, Parkinson & Liddle 2006; Feroz & Hobson 2008; Handley, Hobson & Lasenby 2015), neutron star astrophysics (Meyers et al. 2021a; Meyers et al. 2021b; Melatos et al. 2023), particle physics (Trassinelli 2019), and materials science (Pártay, Bartók & Csányi 2009). For reviews of nested sampling, we refer the reader to Buchner (2023) and Ashton et al. (2022). Multiple nested sampling algorithms and computational libraries exist (e.g. Feroz & Hobson 2008; Feroz, Hobson & Bridges 2009; Handley et al. 2015; Speagle 2020; Buchner 2021). In GW research, it is common to use the DYNESTY sampler (Speagle 2020) via the BILBY (Ashton & Talbot 2021) front-end library. We follow this precedent and use BILBY for all nested sampling Bayesian inference in this work.

The primary tunable parameter in nested sampling is n_{live} . More live points assist with large parameter spaces and multimodal problems, whilst the uncertainties in the evidence and the posterior scale as $\mathcal{O}(n_{\text{live}}^{-1/2})$. However the computational runtime scales as $\mathcal{O}(n_{\text{live}})$. Ashton et al. (2022) offered a rule-of-thumb trade-off, where the minimum number of live points should be greater than the number of static parameters. Informal empirical tests conducted as part of this paper support the trade-off suggested by Ashton et al. (2022); we find typically that the true θ is contained within the 90 per cent credible interval of the one-dimensional marginalized posteriors of $\hat{\theta}$ for $n_{\text{live}} > 7 + 5N$ with $N \leq 50$. Unless stated otherwise, we take $n_{\text{live}} = 1000$ for all results presented in this work. Empirically, we find that values of $n_{\text{live}} > 1000$ do not improve appreciably the accuracy of the results presented in Sections 4 and 5.

3.3 Model selection

The evidence integral \mathcal{Z} returned by nested sampling is the probability of the data \mathbf{Y} given a model \mathcal{M}_i . We compare competing models via a Bayes factor,

$$\beta = \frac{\mathcal{Z}(\mathbf{Y}|\mathcal{M}_1)}{\mathcal{Z}(\mathbf{Y}|\mathcal{M}_0)}. \quad (33)$$

Throughout this work, we take \mathcal{M}_1 to be the state-space model that includes the presence of a GW. \mathcal{M}_0 is the null model, which assumes no GW exists in the data, and is equivalent to setting $g^{(n)}(t) = 1$ in equation (19). The Bayes factors we quote in this work therefore quantify whether the data support evidence for a GW signal compared to there being no GW signal present.

3.4 Summary of workflow

For the reader’s convenience, we now summarize the workflow for a representative PTA analysis using the Kalman filter and nested sampler for parameter estimation and model selection:

- (i) Specify a PTA composed of N pulsars.
- (ii) Obtain N data inputs $f_m^{(n)}(t)$, collectively labelled \mathbf{Y} .
- (iii) Specify a state-space model \mathcal{M} , with static parameters θ .
- (iv) Specify prior distribution $\pi(\theta)$.
- (v) Sample n_{live} points from $\pi(\theta)$.
- (vi) For each live point:
 - (a) Pass the sample θ_{sample} to the Kalman filter.
 - (b) Iterate over the input data using the Kalman filter and obtain a single log \mathcal{L} value, equation (30).

- (vii) Remove the live point with the lowest likelihood value, $\log \mathcal{L}_{\text{lowest}}$.
- (viii) Sample a new live point from $\pi(\theta)$, subject to the requirement that the new likelihood obeys $\mathcal{L}_{\text{new}} > \mathcal{L}_{\text{lowest}}$, where $\log \mathcal{L}_{\text{new}}$ is calculated via steps (6)(1) and (2).
- (ix) Update $p(\theta|Y)$ and \mathcal{Z} with nested sampler.
- (x) Repeat steps (7)–(9) until convergence criteria are satisfied.

In order to compute β the above workflow is repeated for a different \mathcal{M} . The resulting \mathcal{Z} values can then be compared. We remind the reader that the above workflow differs from a realistic PTA analysis in one important respect, namely that the data are input as frequency time-series $f_m^{(n)}(t)$ instead of pulse TOAs. The generalization to TOAs is subtle and will be tackled in a forthcoming paper.

3.5 Relation to traditional PTA analyses

It is natural to ask how the workflow in Sections 3.1–3.4 differs from traditional PTA analyses. One superficial difference is the detection statistic: traditional analyses use a frequentist matched filter like the maximum-likelihood \mathcal{F} -statistic (Ellis et al. 2012; Antoniadis et al. 2023c) or Bayesian inference (Arzoumanian et al. 2023; Agazie et al. 2023c), whereas the workflow in Sections 3.1–3.4 maximizes the Kalman likelihood in equation (30).

A more subtle difference is the manner in which the intrinsic, achromatic timing noise is tracked and modelled. With respect to noise tracking, traditional analyses usually involve some form of least-squares minimization. The adaptive gain, which measures the fractional amount of new information incorporated into the updated state estimate at each time-step, tends to zero in the limit $t \rightarrow \infty$ in least-squares minimization but remains non-zero in the same limit for a Kalman filter. That is, a Kalman filter responds more nimbly to new data than least-squares estimators, which is why it is favoured in many electrical and mechanical engineering applications (Gelb et al. 1974; Zarchan & Musoff 2000; Byrne 2005; Särkkä 2013). The Kalman gain is discussed more fully in Appendix B1 and compared with traditional analyses. With respect to noise modelling, traditional analyses usually introduce a sum of Fourier modes in a TOA-based phase model, with random coefficients drawn from an ensemble-averaged power spectral density, whose form is often a (broken) power law with estimable amplitude and exponent(s). In this paper, the Kalman filter tracks the specific, time-ordered, random realization of the noise in the data consistent with the Ornstein–Uhlenbeck process in equations (2)–(4), without ensemble averaging or imposing the extra time-domain structure implicit in a finite-term Fourier expansion (e.g. continuity from one time-step to the next). Nevertheless, the noise models are similar and can be related in certain limits; for example, equations (2)–(4) imply an associated power spectral density, which is a broken power law. A fuller discussion of equations (2)–(4), including a comparison with traditional noise models, is presented in Appendix B2.

4 VALIDATION WITH SYNTHETIC DATA

In this section, we outline how synthetic data are generated in order to validate the analysis scheme in Section 3. Synthetic data enable validation to occur systematically and under controlled conditions. In Section 4.1, we describe how to construct a representative synthetic PTA, and how to set astronomically reasonable values for the static pulsar parameters θ_{psr} . In Section 4.2, we demonstrate how to solve equations (1)–(4), and (16)–(21) for the synthetic PTA so as to generate noisy, frequency time-series $f_m^{(1)}(t), \dots, f_m^{(N)}(t)$.

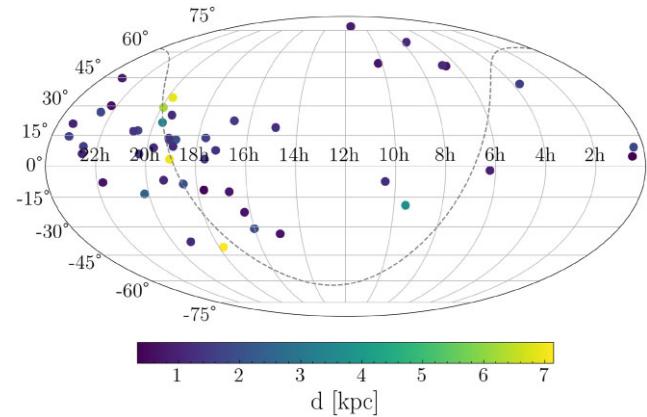


Figure 2. Spatial distribution in Galactic coordinates of 47 pulsars from the 12.5 yr NANOGrav data release that make up the synthetic PTA used in this work. The pulsar distances relative to the observer are also indicated, with distance ≤ 2 kpc for 38 pulsars. The grey dashed curve denotes the Galactic plane.

4.1 Constructing a synthetic PTA

We consider, by way of illustration, a synthetic PTA composed of the 47 pulsars in the 12.5 yr NANOGrav data set (Arzoumanian et al. 2020b). The NANOGrav pulsars are chosen arbitrarily as being representative of a typical PTA; the analysis below extends unchanged to any other PTA.

For each pulsar, we adopt fiducial values for $\mathbf{q}^{(n)}$, $d^{(n)}$, $f_{\text{em}}^{(n)}(t_1)$, and $f_{\text{em}}^{(n)}(t_1)$, with the latter two quantities evaluated at the SSB. A table of fiducial values is presented in Appendix C for the sake of reproducibility. The sky positions and colour-coded distances of the pulsars are displayed in Fig. 2. The pulsar parameters are acquired via the Australia Telescope National Facility (ATNF) pulsar catalogue (Manchester et al. 2005) using the PSRQPY package (Pitkin 2018).

The other static pulsar parameters are $\gamma^{(n)}$ and $\sigma^{(n)}$, for which no direct measurements exist. The ratio $\sigma^{(n)}/[\gamma^{(n)}]^{1/2}$ sets the typical root mean square fluctuations in $f_p^{(n)}(t)$, as discussed in Section 2.1, and the mean reversion time-scale typically satisfies $[\gamma^{(n)}]^{-1} \gg T_{\text{obs}}$ (Price et al. 2012; Meyers et al. 2021a, b; Vargas & Melatos 2023). In this paper, for the sake of simplicity in the absence of independent measurements, we fix $\gamma^{(n)} = 10^{-13} \text{ s}^{-1}$ for all n , consistent with (Vargas & Melatos 2023). We follow two complementary approaches in order to set physically reasonable values for $\sigma^{(n)}$. The first approach relies on the empirical timing noise model from Shannon & Cordes (2010) which gives the standard deviation of the pulsar TOAs, $\sigma_{\text{TOA}}^{(n)}$, as

$$\ln \left[\frac{\sigma_{\text{TOA}}^{(n)}}{\mu \text{s}} \right] = \ln \alpha_1 + \alpha_2 \ln f_p^{(n)} + \alpha_3 \ln \left[\frac{f_p^{(n)}}{10^{-15} \text{s}^{-2}} \right] + \alpha_4 \ln \left(\frac{T_{\text{cad}}}{1 \text{yr}} \right), \quad (34)$$

where T_{cad} is the cadence of the timing observations. For MSPs, the best-fitting parameters are $\ln \alpha_1 = -20 \pm 20$, $\alpha_2 = 1 \pm 2$, $\alpha_3 = 2 \pm 1$, and $\alpha_4 = 2.4 \pm 0.6$ ($\pm 2\sigma$ confidence limits). The uncertainties are broad; for the purpose of generating astrophysically representative synthetic data we adopt the central values in this paper. Throughout this work, we assume for simplicity that all pulsars are observed with a weekly cadence, $T_{\text{cad}} = 1$ week. In order to relate equation (34) to $\sigma^{(n)}$ in equation (1), we equate $\sigma^{(n)}$ to the root mean square TOA noise accumulated over one week, obtaining

$$\sigma^{(n)} \approx \sigma_{\text{TOA}}^{(n)} f_p^{(n)} T_{\text{cad}}^{-3/2}. \quad (35)$$

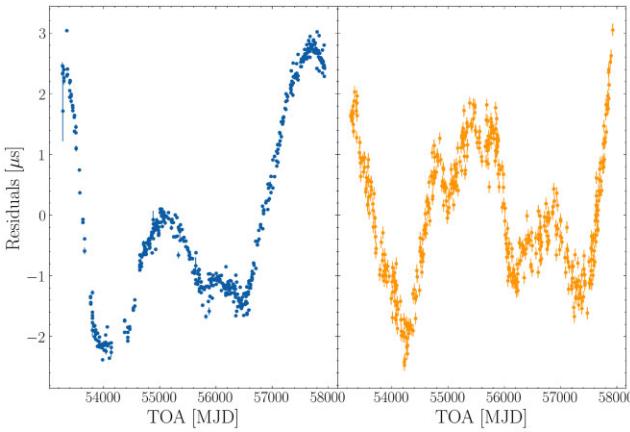


Figure 3. Actual (left panel, blue points) and synthetic (right panel, orange points) phase residuals for NANOGrav MSP PSR J1939 + 2134. The actual residuals are sourced from the NANOGrav 12.5 yr wide-band timing data set (Pennucci & Collaboration 2020; Alam et al. 2021). The synthetic residuals are generated by numerically solving equation (1) with $\gamma^{(n)} = 10^{-13} \text{ s}^{-1}$ and $\sigma^{(n)} = 2 \times 10^{-27} \text{ s}^{-3/2}$. The error bars indicate the uncertainty in the phase residuals and are generated by propagating the uncertainty in the TOAs through TEMPO2. We set the TOA uncertainty to be a constant, viz. $0.1\mu\text{s}$. The blue and orange residuals are qualitatively similar by inspection. Similar results are obtained for the other 46 pulsars in the synthetic PTA in Section 4.1 by tuning the value of $\sigma^{(n)}$.

For the synthetic NANOGrav PTA depicted in Fig. 2, the median $\sigma^{(n)}$ calculated in this way is $\sigma^{(n)} = 5.51 \times 10^{-24} \text{ s}^{-3/2}$, with $\min[\sigma^{(n)}] = 1.67 \times 10^{-26} \text{ s}^{-3/2}$ for PSR J0645 + 5158 and $\max[\sigma^{(n)}] = 2.56 \times 10^{-19} \text{ s}^{-3/2}$ for PSR J1939 + 2134.

As a cross-check, we estimate $\sigma^{(n)}$ by solving equation (1) numerically using the BABOO package,² generating a synthetic phase solution,

$$\varphi^{(n)}(t) = \int_0^t dt' f_p^{(n)}(t') , \quad (36)$$

and adjusting $\sigma^{(n)}$ iteratively to generate phase residuals which qualitatively (i.e. visually) resemble empirical phase residuals measured from real pulsars; see section 2 in Vargas & Melatos (2023) for a successful example of this approach. We obtain empirical phase residuals from the NANOGrav 12.5 yr wide-band timing data set (Pennucci & Collaboration 2020; Alam et al. 2021). A visual cross-check is sufficient for the purposes of this paper, where we seek broadly representative values for $\sigma^{(n)}$. In an analysis involving real PTA data, $\sigma^{(n)}$ would be estimated from the data jointly with the other static parameters in θ . In Fig. 3, we compare the synthetic and empirical residuals for PSR J1939 + 2134, one of the 47 NANOGrav pulsars plotted in Fig. 2. We see that the synthetic and empirical residuals are qualitatively similar. For this pulsar, the synthetic residuals are generated using $\sigma^{(n)} = 2 \times 10^{-27} \text{ s}^{-3/2}$ which is much less than the central value $\sigma^{(n)} = 2.56 \times 10^{-19} \text{ s}^{-3/2}$ inferred from equations (34) and (35), but is consistent with the wide ranges on $\alpha_1, \dots, \alpha_4$, for example, $\ln \alpha_1 = -38$ implies $\sigma^{(n)} = 3.90 \times 10^{-27} \text{ s}^{-3/2}$ for PSR J1939+2134.

4.2 Generating a synthetic sequence of pulse frequencies

We generate N synthetic noisy time series of the measured pulse frequency $f_m^{(n)}(t)$, one for each pulsar $1 \leq n \leq N$, as follows:

- (i) Integrate equations (1)–(4) numerically for the synthetic PTA described in Section 4.1, to obtain random realizations of $f_p^{(n)}(t)$ for $1 \leq n \leq N$.
- (ii) Map from $f_p^{(n)}(t)$ to $f_m^{(n)}(t)$ via equations (16) and (21).

Equations (1)–(4) are solved by a Runge–Kutta Itô integrator implemented in the SDEINTPYTHON package.³ The static pulsar parameters θ_{psr} are completely specified for the synthetic PTA outlined in Section 4.1. For this work, we consider all pulsars to be observed for $T_{\text{obs}} = 10 \text{ yr}$, uniformly sampled with a weekly cadence.

The measurement noise covariance as defined in equation (18) can be approximately related to the uncertainty in the pulse TOA, σ_{TOA} , through

$$\sigma_m \approx \sigma_{\text{TOA}} f_p^{(n)} T_{\text{cad}}^{-1} . \quad (37)$$

Although equations (35) and (37) superficially resemble one another, they are distinct. Equation (35) deals with the timing noise intrinsic to the pulsar due to rotational irregularities, whereas equation (37) handles the detector measurement noise. For an MSP with $f_p^{(n)} \sim 10^2 \text{ Hz}$ observed with weekly cadence and $\sigma_{\text{TOA}} \sim 1\mu\text{s}$, equation (37) implies $\sigma_m \sim 10^{-10} \text{ Hz}$. The most accurately timed pulsars have $\sigma_{\text{TOA}} \sim 10 \text{ ns}$ and $\sigma_m \sim 10^{-12} \text{ Hz}$. Throughout this paper, we fix $\sigma_m = 10^{-11} \text{ Hz}$ and take it as known a priori rather than a parameter to be inferred for the sake of simplicity; this assumption can be relaxed easily when analysing real data. Whilst σ_m is the same for every pulsar, $f_m^{(n)}(t)$ is constructed from a different random realization of $\varepsilon^{(n)}(t)$ for each pulsar.

Finite precision arithmetic leads to numerical errors when solving equations (1)–(4) in the regime $|\sigma^{(n)} dB(t)| \ll f_p^{(n)}$ relevant to PTAs, where $dB(t)$ labels an increment of Brownian motion (see Appendix A). To fix the problem, we subtract the deterministic frequency evolution and track the new variable

$$f_p^{*(n)} = f_p^{(n)} - f_{\text{em}}^{(n)} , \quad (38)$$

equivalent to a change of variables. We similarly modify the measurement variable to be

$$f_m^{*(n)} = f_m^{(n)} - f_{\text{em}}^{\diamond(n)} , \quad (39)$$

where $f_{\text{em}}^{\diamond(n)}$ is a guess of the deterministic evolution based on the pulsar ephemeris returned by TEMPO2 which is measured to high accuracy in practice. For synthetic data, we can set $f_{\text{em}}^{\diamond(n)} = f_{\text{em}}^{(n)}$ without loss of generality, but this is impossible generally for astronomical observations, because the spin-down ephemeris is only known approximately. Equation (16) is then updated to read

$$f_m^{*(n)}(t) = f_p^{*(n)}(t - d) g^{(n)}(t) - f_{\text{em}}^{(n)}(t - d) [1 - g^{(n)}(t)] . \quad (40)$$

We emphasize that the change of variables in equations (38) and (39) is a convenient device to bring the numerical values into a reasonable dynamic range without having to use excessively long floating point formats (e.g. long double and quadruple). It does not remove any degrees of freedom nor does it involve an approximation. In particular $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$ remain as static parameters but appear in the measurement equation (40) rather than the dynamical state equations, that is, equations (1)–(4).

²<https://github.com/meyers-academic/baboo>

³<https://github.com/mattja/sdeint>

Table 1. Summary of injected static parameters used for generating synthetic data in the representative example of Section 5, along with the prior used for Bayesian inference on each parameter (rightmost column). The top and bottom halves of the table contain the values for θ_{gw} and θ_{psr} , respectively. The subscript ‘ATNF’ denotes values obtained from the ATNF pulsar catalogue as described in Section 4.1. The subscript ‘SC’ indicates that the injected value is calculated using equations (34) and (35) in Shannon & Cordes (2010). The quantities $\eta_f^{(n)}$ and $\eta_{\dot{f}}^{(n)}$ are the errors in $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$ respectively, as quoted in the ATNF catalogue. No priors are set for $d^{(n)}$ and $\gamma^{(n)}$, as those parameters do not enter the inference model, as discussed in Section 5.1, cf. equation (41).

	Parameter	Injected value	Units	Prior
θ_{gw}	Ω	5×10^{-7}	Hz	LogUniform(10^{-9} , 10^{-5})
	α	1.0	rad	Uniform(0, 2π)
	δ	1.0	rad	Cosine($-\pi/2$, $\pi/2$)
	ψ	0.90	rad	Uniform(0, $\pi/2$)
	Φ_0	3.30	rad	Uniform(0, 2π)
	h_0	5×10^{-15}	—	LogUniform(10^{-15} , 10^{-9})
	ι	1.0	rad	Sin(0, π)
θ_{psr}	$f_{\text{em}}^{(n)}(t_1)$	$f_{\text{ATNF}}^{(n)}$	Hz	Uniform $\left[f_{\text{ATNF}}^{(n)} - 10^3 \eta_f^{(n)}, f_{\text{ATNF}}^{(n)} + 10^3 \eta_f^{(n)} \right]$
	$\dot{f}_{\text{em}}^{(n)}(t_1)$	$\dot{f}_{\text{ATNF}}^{(n)}$	s^{-2}	Uniform $\left[\dot{f}_{\text{ATNF}}^{(n)} - 10^3 \eta_{\dot{f}}^{(n)}, \dot{f}_{\text{ATNF}}^{(n)} + 10^3 \eta_{\dot{f}}^{(n)} \right]$
	$d^{(n)}$	$d_{\text{ATNF}}^{(n)}$	m	—
	$\sigma^{(n)}$	$\sigma_{\text{SC}}^{(n)}$	$\text{s}^{-3/2}$	LogUniform(10^{-26} , 10^{-22})
	$\gamma^{(n)}$	10^{-13}	s^{-1}	—

5 REPRESENTATIVE END-TO-END ANALYSIS

In this section, we apply the analysis scheme in Section 3 and the validation procedure in Section 4 to a PTA perturbed by a GW from an individual quasi-monochromatic SMBHB source. The analysis of a stochastic background composed of the superposition of multiple sources is more challenging and is postponed to a forthcoming paper. The goal of this section is to run end-to-end through every step in the analysis for a representative worked example to help the reader implement the scheme and reproduce the results in Sections 6 and 7. To assist with reproducibility, we present intermediate outputs from each step. In Section 5.1, we define the priors on θ . In Section 5.2, we define the measurement equation used in the inference model by the Kalman filter. In Section 5.3, we apply the workflow in Section 3.4 to estimate θ , for a single realization of the pulsar process noise and the measurement noise. In Section 5.4, we extend the parameter estimation exercise to multiple noise realizations to quantify the variance in the parameter estimates. In Section 5.5, we calculate the detectability of the source as a function of h_0 .

The static GW source parameters θ_{gw} used for this mock analysis are selected to be astrophysically reasonable and representative. The injected components of θ_{gw} and θ_{psr} are summarized in the second column of Table 1.

5.1 Prior distribution

The first step is to select a reasonable Bayesian prior, $\pi(\theta)$, for the static parameters. For $\pi(\theta_{\text{gw}})$, we choose standard non-informative priors (e.g. Bhagwat et al. 2021; Falxa et al. 2023) as summarized in Table 1. The static parameters Φ_0 and ψ are degenerate; a GW with (Φ_0, ψ) is identical to a GW with $(\Phi_0 + \pi, \psi + \pi/2)$. The degeneracy is well known in the PTA literature (Bécsy, Cornish & Digman 2022; Charisi et al. 2024) and results in bimodal posteriors for Φ_0 and ψ . For the purposes of this paper, it is sufficient to circumvent the issue by restricting the prior on ψ to the domain $0 \leq \psi \leq \pi/2$. A similar approach is taken in targeted searches of continuous GWs by the Laser Interferometer Gravitational-Wave Observatory (LIGO; Prix & Krishnan 2009). In the hypothetical event that we do not restrict the prior we obtain bimodal posteriors in

Φ_0 and ψ (cf. Section 5.3) which are out of phase by $(\Phi_0 + \pi, \psi + \pi/2)$, as expected, in line with, for example, fig. 2 of Bécsy et al. (2022). The degeneracy between Φ_0 and ψ is related to the formal concept of identifiability from the theory of signal processing, which is applied routinely to engineering problems (Bellman & Åström 1970). Identifiability refers to whether it is theoretically possible to infer unique and accurate parameter values, given the measured data and the model structure (Walter & Pronzato 1996; Dobre et al. 2012; Guillaume et al. 2019; Casella & Berger 2021).

We now discuss the choice of $\pi(\theta_{\text{psr}})$. The parameters that govern the deterministic evolution of $f_{\text{p}}^{(n)}(t)$, namely $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$, are well determined by radio timing observations. We identify $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$ with the pulsar barycentric rotation frequency and its time derivative respectively, as quoted in catalogues (Manchester et al. 2005). For the 12.5 yr NANOGrav pulsars, the median fractional errors on $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$ are $\pm 2.68 \times 10^{-13}$ per cent and $\pm 2.31 \times 10^{-3}$ per cent respectively. In this paper, we adopt uniform priors on $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$, which extend $\pm 10^3 \eta_f^{(n)}$ and $\pm 10^3 \eta_{\dot{f}}^{(n)}$, respectively about the central, injected values, where $\eta_f^{(n)}$ and $\eta_{\dot{f}}^{(n)}$ denote the errors quoted in the ATNF Pulsar Database (Manchester et al. 2005). Wider-than-necessary priors, such as those above, test the method more stringently than narrow priors. The results below confirm that the method estimates θ_{gw} accurately, whether the priors are wide or narrow.

The pulsar distances $d^{(n)}$ are less constrained than $f_{\text{em}}^{(n)}(t_1)$ and $\dot{f}_{\text{em}}^{(n)}(t_1)$, with typical uncertainties ~ 10 per cent (Yao et al. 2017; Arzoumanian et al. 2018a). In this paper, however, we omit the term involving $d^{(n)}$ from equation (21), as justified in Section 5.3, so there is no need to set a prior on $d^{(n)}$. Furthermore we do not set a prior on $\gamma^{(n)}$, because we have $\gamma^{(n)} T_{\text{obs}} \sim 10^{-5}$, and $\gamma^{(n)}$ is effectively ‘unobservable’ over a decade; that is, for $T_{\text{obs}} = 10$ yr, the solution of equation (1) is approximately independent of $\gamma^{(n)}$, as long as $[\gamma^{(n)}]^{-1} \gg T_{\text{obs}}$ is satisfied. It is therefore sufficient for validation purposes to carry $\gamma^{(n)}$ through the analysis at its injected value.⁴

⁴We check this empirically in two independent ways for an informal selection of test cases: (i) we set an uninformative prior $\pi[\gamma^{(n)}/1 \text{ s}^{-1}] \sim$

Most pulsars in the synthetic PTA have $10^{-25} \leq \sigma^{(n)}/(1 \text{ s}^{-3/2}) \leq 10^{-23}$ as calculated from equations (34) and (35). For these pulsars, we set an uninformative broad prior $\pi[\sigma^{(n)}/(1 \text{ s}^{-3/2})] \sim \text{LogUniform}(10^{-26}, 10^{-22})$. The sole exception is PSR J1939 + 2134 which has $\sigma^{(n)} \sim 10^{-19} \text{ s}^{-3/2}$. For validation purposes, we artificially set $\sigma^{(n)} = 10^{-23} \text{ s}^{-3/2}$ for this pulsar, so that $\sigma^{(n)}$ for every pulsar in the synthetic PTA falls within the aforementioned log-uniform prior. This is done purely for testing convenience; it is straightforward to expand the prior when analysing real, astronomical data.

By not setting priors on $\gamma^{(n)}$ and $d^{(n)}$, we reduce the number of parameters to $7 + 3N$. We use the notation $\boldsymbol{\theta}_{\text{psr, reduced}}$ to refer to the reduced parameter set, cf. equation (25). Explicitly, we write

$$\boldsymbol{\theta}_{\text{psr, reduced}} = \{f_{\text{em}}^{(n)}(t_1), \dot{f}_{\text{em}}^{(n)}(t_1), \sigma^{(n)}\}_{1 \leq n \leq N}. \quad (41)$$

The injected static parameters and their corresponding priors for the representative analysis in Section 5 are summarized in Table 1.

5.2 Earth and pulsar terms

The general measurement equation used in the inference model, equation (21), separates into four terms. The first two terms, $H_{ij}^{(+)} \cos \Phi(t) + H_{ij}^{(\times)} \sin \Phi(t)$, depend only on the GW source parameters which are shared across all pulsars. The argument of the trigonometric functions corresponds to the GW phase at the Earth. The next two terms, $H_{ij}^{(+)} \cos \Phi^{(n)}(t) + H_{ij}^{(\times)} \sin \Phi^{(n)}(t)$, depend additionally on $d^{(n)}$ and $\mathbf{q}^{(n)}$ and vary between pulsars. The argument of the trigonometric functions corresponds to the GW phase at each individual pulsar. The first and second pairs of terms are commonly referred to as the ‘Earth term’ and ‘pulsar term’, respectively. Whilst the Earth term is phase coherent between all pulsars, the pulsar terms have uncorrelated phases. They are typically considered as a source of self-noise and dropped from many standard PTA analyses (e.g. Sesana & Vecchio 2010; Babak & Sesana 2012; Petiteau et al. 2013; Zhu et al. 2015b; Taylor et al. 2016; Goldstein et al. 2018; Charisi et al. 2024) at the expense of a modestly reduced detection probability (~ 5 per cent) and the introduction of a bias in the inferred sky position (Zhu et al. 2016; Chen & Wang 2022).

In this paper, we follow the standard approach and drop the pulsar terms from the inference calculation (but not the source model). Explicitly the measurement equation used in the Kalman filter reduces to

$$f_{\text{m}}^{(n)}(t) = f_{\text{p}}^{(n)}(t - d)g_{\text{Earth}}^{(n)}(t), \quad (42)$$

with

$$g_{\text{Earth}}^{(n)}(t) = 1 - z_{\text{Earth}}^{(n)}(t), \quad (43)$$

and

$$z_{\text{Earth}}^{(n)}(t) = \frac{[q^{(n)}]^i [q^{(n)}]^j}{2[1 + \mathbf{n} \cdot \mathbf{q}^{(n)}]} \left[H_{ij}^{(+)} \cos \Phi(t) + H_{ij}^{(\times)} \sin \Phi(t) \right]. \quad (44)$$

We defer the inclusion of the pulsar terms in the inference calculation to a future paper. As discussed in Section 5.1, this choice reduces the dimensionality of the parameter estimation problem because the measurement equation, equation (43), is no longer a function of the pulsar distance. We stress that the pulsar terms are dropped only when doing Bayesian inference, that is, from the Kalman filter model that

$\text{LogUniform}(10^{-15}, 10^{-10})$; and (ii) we deliberately displace $\gamma^{(n)}$ in the inference analysis from its true, injected value in the synthetic data, for example, $\gamma^{(n)} = 10^{-14} \text{ s}^{-1}$ versus 10^{-13} s^{-1} , respectively. The results from (i) and (ii) are found to be the same as those reported in Section 5.

feeds into the nested sampling algorithm. The synthetic data include the pulsar terms in full.

5.3 Posterior distribution

In this section, we calculate the posterior probability distribution for the static parameters $\boldsymbol{\theta} = \boldsymbol{\theta}_{\text{gw}} \cup \boldsymbol{\theta}_{\text{psr}}$ and compare it to the known, injected values. The aims are (i) to demonstrate that the analysis scheme works (i.e. that it converges to a well-behaved, unimodal posterior), and (ii) to give a preliminary sense of its accuracy. Initially we consider a single noise realization when generating the synthetic data for the representative example in Table 1. We apply the Kalman filter in conjunction with nested sampling in order to infer the joint posterior distribution $p(\boldsymbol{\theta} | \mathbf{Y})$.

Fig. 4 displays results for the seven parameters in $\boldsymbol{\theta}_{\text{gw}}$ in the form of a traditional corner plot. The histograms are the one-dimensional posteriors for each parameter, marginalized over the six other parameters. The dashed vertical blue lines mark the 0.16 and 0.84 quantiles; the solid orange line marks the known injected value. The two-dimensional contours mark the $(0.5, 1, 1.5, 2)\sigma$ level surfaces. All histograms and contours are consistent with a unimodal joint posterior, which peaks near the known, injected values. There is scant evidence of railing against the prior bounds. For this representative example, with characteristic strain $h_0 = 5 \times 10^{-15}$, the analysis scheme estimates $\boldsymbol{\theta}_{\text{gw}}$ accurately. The injected values are contained within the 90 per cent credible interval for five of the seven parameters in $\boldsymbol{\theta}_{\text{gw}}$.

The posterior is approximately symmetric about the known injected value for some parameters such as Ω , where the posterior median and the injected value coincide approximately. For other parameters (e.g. ψ , ι) the distribution is not symmetric about the injected value and the posterior median and the injected value do not coincide. The median value of the posterior for ι is shifted by ≈ 0.13 rad relative to the injected value, although the injected value does still remain inside the 90 per cent credible interval. Similar effects are seen, albeit with a smaller shift, in other parameters such as δ and α . For a single realization of the noise it is not clear if this discrepancy is a systematic effect, that is, a bias, or a random outcome specific to this particular noise realization. We explore this further in Sections 5.4 and 7 using more noise realizations and show that indeed there is a systematic bias from dropping the pulsar terms, similar to that reported by Zhu et al. (2016).

Similar results are obtained for the $3N$ parameters in $\boldsymbol{\theta}_{\text{psr, reduced}}$. Again, the parameters are recovered unambiguously, in the sense that the nested sampler converges smoothly to a unimodal joint posterior near the known, injected parameter values, with all $3N$ injected parameters lying within the 90 per cent credible interval of the $3N$ one-dimensional, marginalized posteriors. We do not display the resulting corner plot because it is too big ($3N = 141$), and because inferring $\boldsymbol{\theta}_{\text{gw}}$ is the focus of this paper and contemporary PTA analyses.

5.4 Multiple noise realizations: dispersion of outcomes

The results in Section 5.3 are obtained for a single realization of the noise processes $\xi^{(n)}(t)$ and $\varepsilon^{(n)}(t)$. It is important to confirm that the analysis scheme returns accurate answers for arbitrary noise realizations and that the specific realization of the noisy data used in Section 5.3 is not particularly advantageous by accident. It is also important to quantify, albeit approximately, the natural random dispersion in the one-dimensional posterior medians from one noise realization to the next, as the dispersion is a practical measure of the

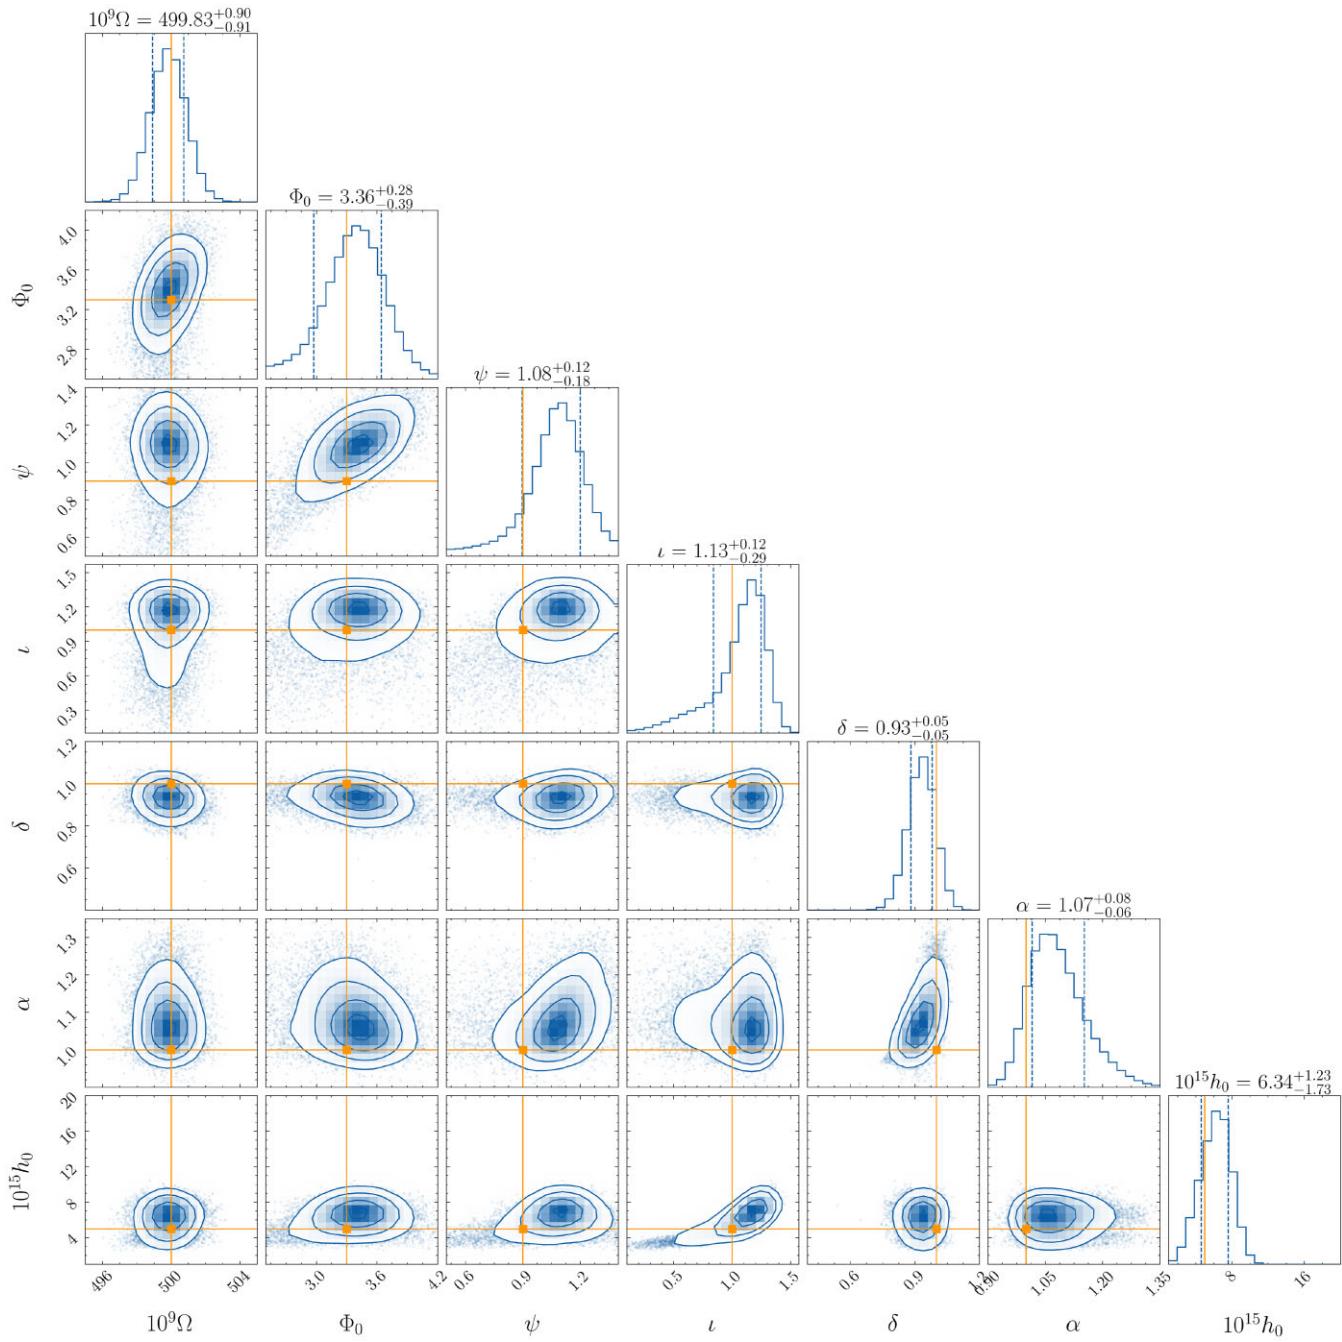


Figure 4. Posterior distribution of the GW source parameters θ_{gw} for the representative system in Table 1, for a single realization of the system noise. The horizontal and vertical orange lines indicate the true injected values. The contours in the two-dimensional histograms mark the $(0.5, 1, 1.5, 2)\sigma$ levels after marginalizing over all but two parameters. The one-dimensional histograms correspond to the joint posterior distribution marginalized over all but one parameter. The super titles of the marginalized one-dimensional histograms specify the posterior median and the 0.16 and 0.84 quantiles. We plot the scaled variables $10^9\Omega$ (units: rad s^{-1}) and $10^{15}h_0$. The Kalman filter and nested sampler estimate accurately all seven parameters in θ_{gw} , although biases are observed in some parameters. The horizontal axes span a subset of the prior domain for all seven parameters.

accuracy of the parameter estimation scheme, when it is applied to real astronomical data, where the true parameter values and specific noise realization are unknown.

To this end, we start with the representative example in Table 1 and generate 1000 realizations of the process noise $\xi^{(n)}(t)$ and measurement noise $\varepsilon^{(n)}(t)$. For each realization, we independently estimate the static parameters, θ . In Fig. 5, we plot the estimates of θ_{gw} for 100 arbitrary realizations. The seven parameters in θ_{gw} are

recovered unambiguously. The corner plot is arranged identically to Fig. 4, which stems from one realization. We plot 100 realizations rather than the full set of 1000 to avoid overcrowding. As in Section 5.3, θ_{psr} is also recovered unambiguously across the 1000 noise realizations, but for the sake of brevity we do not show the results here.

Fig. 5 confirms two main points: (i) the results from the 100 noise realizations overlap with the single realization in Fig. 4; and (ii)

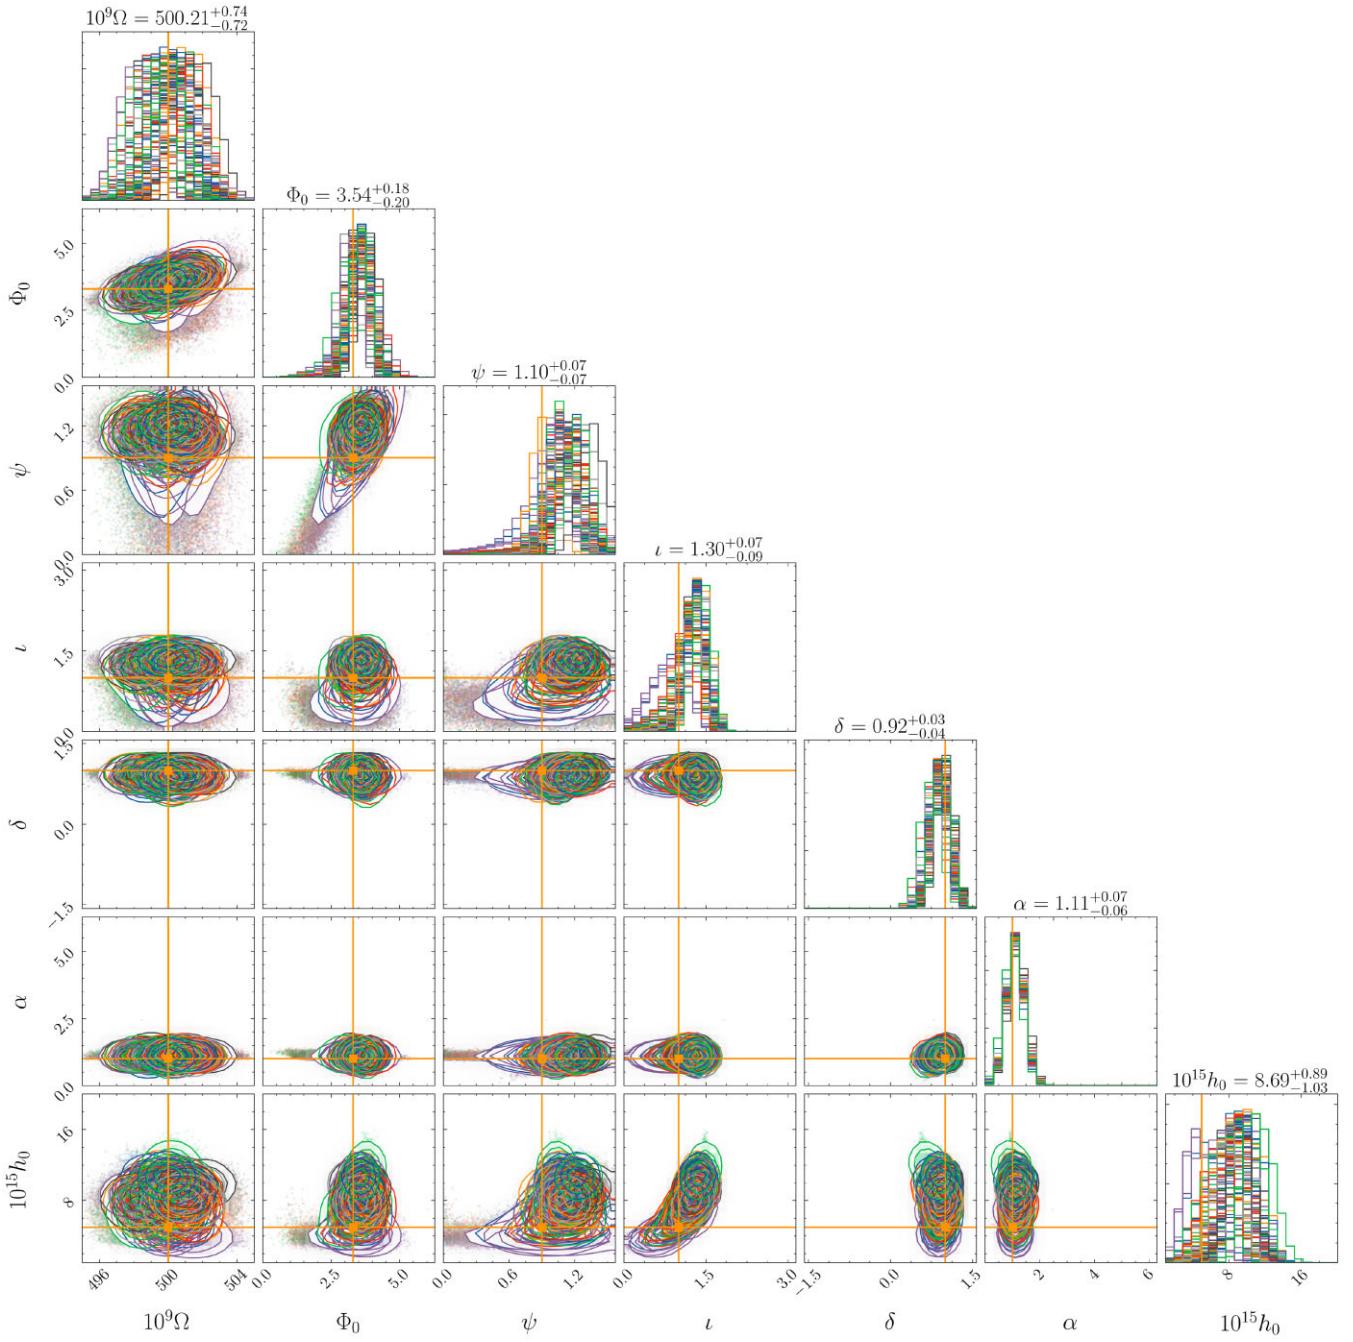


Figure 5. Same as Fig. 4, but for 100 realizations of the noise processes, with curves coloured differently. The super titles of the one-dimensional histograms record the posterior median and the 0.16 and 0.84 quantiles of the median realization. The known, injected value lies within the 90 per cent credible interval for 597 out of the 700 combinations of seven parameters and 100 noise realizations. There is an appreciable dispersion among the peaks of the one-dimensional posteriors, with coefficient of variation ~ 10 per cent across the 700 combinations. A slight skew bias is apparent in some of the parameters, for example, ι .

the dispersion among the peaks of the one-dimensional posteriors is appreciable, with variations of ~ 10 per cent (coefficient of variation) across the seven parameters and 100 realizations. Indeed, considering the one-dimensional marginalized posteriors, we find that the injected value is contained within the 90 per cent credible interval in 597 (i.e. 85 per cent) out of the 700 possible combinations of the seven parameters and 100 realizations. Fig. 5, like Fig. 4, displays tentative signs of bias, where the one-dimensional posteriors are not symmetric about the injected value. For example, the maximum a posteriori probability estimates of ι appear right-skewed,

consistently overestimating the injected value in all 100 realizations. Similar trends are seen in ψ and h_0 . However, the width of the posteriors is comparable to the putative bias, so it is difficult to draw strong conclusions. Biases are discussed in detail in Section 7. There is no strong evidence for correlations between parameter pairs, for example, banana-shaped contours, except arguably $\iota-h_0$.

We now consider the complete set of 10^3 noise realizations, going beyond the subset of nine realizations in the preceding discussion. We ask the question: how ‘similar’ are the 10^3 marginalized, one-dimensional posteriors computed for each of the seven parameters in

Table 2. Median value of the WD $W_{1,\text{median}}$, for each parameter in θ_{gw} , calculated across the $\binom{10^3}{2}$ pairs of probability posteriors, for the 10^3 noise realizations in Fig. D1. $W_{1,\text{median,inj}}$ is the $W_{1,\text{median}}$ value normalized by the injection value. $W_{1,\text{median,prior}}$ is the $W_{1,\text{median}}$ value normalized by the width of the prior domain (cf. Table 1). Both $W_{1,\text{median,inj}}$ and $W_{1,\text{median,prior}}$ are quoted as percentages. $W_{1,\text{median}}$ is generally smaller than the scales set by the injection values and the prior domain.

Parameter	Injected value	Units	$W_{1,\text{median}}$	$W_{1,\text{median,inj}}$ (per cent)	$W_{1,\text{median,prior}}$ (per cent)
Ω	5×10^{-7}	Hz	8.6×10^{-10}	0.2	8.6×10^{-3}
Φ_0	3.3	rad	0.20	6.2	3.1
ψ	0.9	rad	0.09	9.9	2.8
ι	1.0	rad	0.11	10.5	3.3
δ	1.0	rad	0.04	4.0	1.3
α	1.0	rad	0.06	5.5	0.9
h_0	5×10^{-15}	—	1.3×10^{-15}	25.0	1.3×10^{-4}

θ_{gw} ? There is no unique way to answer this question. In this paper, we appeal to the Wasserstein distance (WD, Vaserstein 1969; Villani 2009) from optimal transport theory, which defines an intuitive notion of similarity between probability distributions. The WD is a popular metric in machine learning (Arjovsky, Chintala & Bottou 2017), climate modelling (Paxton et al. 2022; Kimpson et al. 2023), computational biology (González-Delgado et al. 2023), and geophysics (Morris, Lipp & Roberts 2023); a short review is presented in Appendix D. The WD measures the cost of an optimal strategy for moving probability mass between two distributions from position x to position y , with respect to some cost function $c(x, y)$. In this paper, we use the first WD moment, $W_1(\mu, \nu)$, which is defined and interpreted in Appendix D. For our purposes, W_1 has a convenient physical interpretation given by the Monge–Kantorovich duality (Villani 2003, 2009), viz.

$$|\mathbf{E}(X_\mu) - \mathbf{E}(Y_\nu)| \leq W_1(\mu, \nu), \quad (45)$$

where X_μ and Y_ν are random variates drawn from the distributions μ and ν respectively, and \mathbf{E} denotes the expected value. That is, W_1 bounds the difference in the expectation value of a parameter selected from θ_{gw} with respect to the probability distributions μ and ν . Taking a concrete example, suppose that we infer two one-dimensional posterior distributions $\mu(\iota)$ and $\nu(\iota)$ for ι , for two different realizations of the noise, and calculate $W_1(\mu, \nu) = 0.5$ rad. Then, we can conclude that $|\mathbf{E}(\iota)_\mu - \mathbf{E}(\iota)_\nu| \leq 0.5$ rad.

Table 2 summarizes the WD between the 5×10^5 pairs of one-dimensional posteriors across the 10^3 realizations, for each of the seven parameters in θ_{gw} . The median W_1 for each parameter is tabulated in the fourth column and denoted by $W_{1,\text{median}}$. The $W_{1,\text{median}}$ normalized by the injected, known, value is denoted by $W_{1,\text{median,inj}}$ and quoted in the fifth column as a percentage. $W_{1,\text{median,inj}}$ ranges from 0.17 per cent for Ω to 25 per cent for h_0 . These values quantify approximately the natural dispersion in parameter estimates when the analysis scheme is applied to real, astronomical data, where the true parameter values are unknown. The $W_{1,\text{median}}$ normalized by the width of the prior domain is denoted by $W_{1,\text{median,prior}}$ and quoted in the sixth column as a percentage. $W_{1,\text{median,prior}}$ ranges from 1.3×10^{-4} per cent for h_0 to 3.3 per cent for ι . These values confirm that the nested sampler converges reliably to a single, narrow peak without railing against the prior bounds for 10^3 noise realizations. The $W_{1,\text{median}}$ normalized by the injected, known, value is denoted by $W_{1,\text{median,inj}}$ and quoted in the fifth column as a percentage. $W_{1,\text{median,inj}}$ ranges from 0.17 per cent for Ω to 25 per cent for h_0 .

Further analysis of the WD is performed in Appendix D where we use it for two separate purposes: (i) to measure the similarity

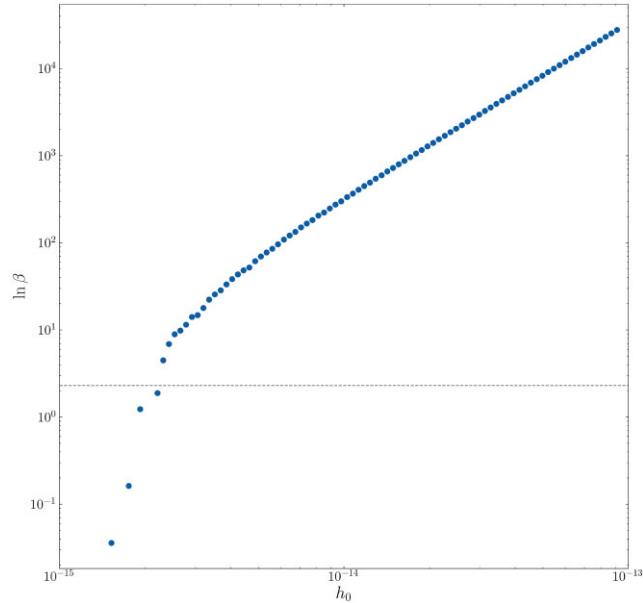


Figure 6. Log Bayes factor (odds ratio) $\ln \beta$ between the competing models \mathcal{M}_1 (GW present in data) and \mathcal{M}_0 (GW not present in data) at different GW strains, h_0 , for the representative example in Table 1. The horizontal grey dashed line labels an arbitrary detection threshold, $\beta = 10$. The minimum detectable strain, for $\beta < 10$, equals 3×10^{-15} . Missing points for $\beta \lesssim 10$ occur when noise subverts the hierarchical relationship between \mathcal{M}_0 and \mathcal{M}_1 (see Section 5.5). Note that the vertical axis features a logarithmic scale for $\ln \beta$ not β . This is appropriate, because β grows exponentially with h_0 .

between probability distributions and (ii) as a convenient heuristic for assessing convergence of nested sampling.

5.5 Detectability versus h_0

We frame the problem of detecting a GW in noisy PTA data in terms of the Bayesian model selection procedure described in Section 3.3. In equation (33), \mathcal{M}_1 is the Earth-term-only model with a GW present, that is, the state-space model with a Kalman filter based on equation (42). The Bayes factor, β , defined in equation (33) is plotted logarithmically in Fig. 6 for the representative source in Table 1, except that we now vary the source amplitude, h_0 , from 10^{-15} (undetectable) to 10^{-12} (easily detectable). To control the test, the noise processes in the synthetic data are identical realizations for

each value of h_0 ; the only change from one h_0 value to the next is h_0 itself.⁵

We see in Fig. 6 an approximate quadratic relationship $\ln \beta \propto h_0^2$ for $h_0 \gtrsim 10^{-14}$. The GW source is detectable with decisive evidence ($\beta \geq 10$) for $h_0 \gtrsim 2.2 \times 10^{-15}$. Of course, the minimum detectable strain is particular to the system in Table 1. It is influenced in general by T_{obs} , θ_{gw} , and θ_{psr} , as discussed in Section 6. Adjusting σ_m we find that an approximate quadratic relationship $\ln \beta \propto h_0^2/\sigma_m^2$ also exists, where h_0/σ_m is the effective signal-to-noise ratio (SNR).

The log odds ratio $\ln \beta$ drops off (i.e. β levels off at unity) for $h_0 \lesssim 3 \times 10^{-15}$. This happens because the two competing models becoming increasingly indistinguishable once the measurement noise dominates the GW signal. Moreover, the points in Fig. 6 become sparser for $h_0 \lesssim 3 \times 10^{-15}$. This happens because we obtain $\ln \beta < 0$ for the missing, intermediate points. This is a noise artefact of the nested sampler. When the sampler converges suboptimally, the hierarchical relationship between \mathcal{M}_0 and \mathcal{M}_1 fails, and $\mathcal{Z}(\mathbf{Y}|\mathcal{M}_1) > \mathcal{Z}(\mathbf{Y}|\mathcal{M}_0)$ no longer holds. There is nothing special about the particular missing points; if one recreates the $\beta(h_0)$ curve by rerunning the nested sampler with another random seed, a different set of points are missing.

6 SMBHB SOURCE PARAMETERS

Section 5 focuses on a single representative system, parametrized in Table 1. In this section, we test the method for various systems, varying the SMBHB source parameters through astrophysically relevant ranges.

We analyse 200 injections constructed by fixing $h_0 = 5 \times 10^{-15}$ and $\iota = 1.0$ rad and drawing the remaining five elements of θ_{gw} randomly from the prior distributions defined in Table 1. We fix h_0 and ι in order to maintain an approximately constant SNR across the 200 injections. For each injection, we compute the posterior distribution of θ_{gw} . To summarize the results we use a parameter-parameter (PP) plot (Cook, Gelman & Rubin 2006). A PP plot displays the fraction of injections included within a given credible interval of the estimated posterior, plotted as a function of the credible interval itself. In the ideal case of perfect recovery, the PP plot should be a diagonal line of unit slope.

Fig. 7 displays the results of the numerical experiment described in the previous paragraph. The shaded grey contours enclose the 1σ , 2σ , and 3σ significance levels for 200 injections. We see that only Ω falls wholly within the 3σ shaded region. The PP curves for the other parameters deviate from the diagonal of unit slope. The deviation is more pronounced for α and δ and less for ψ and Φ_0 . The shape of the graph indicates that the posteriors for these parameters are overconstrained; there are fewer injections contained within higher value credible intervals than would be expected statistically, and there are more injections contained within lower value credible intervals. This stems from the bias noted in Section 5.4. The origin of the bias is discussed in Section 7.

7 SYSTEMATIC BIASES

The tests in Sections 5 and 6 present preliminary evidence that the inferred value of θ_{gw} (e.g. the medians of the one-dimensional marginalized posteriors) is biased away from the true injected value. The biases result from dropping the pulsar terms from the Kalman

⁵Changing the noise realizations as well, from one value of h_0 to the next, adds uninformative scatter to the trend in Fig. 6.

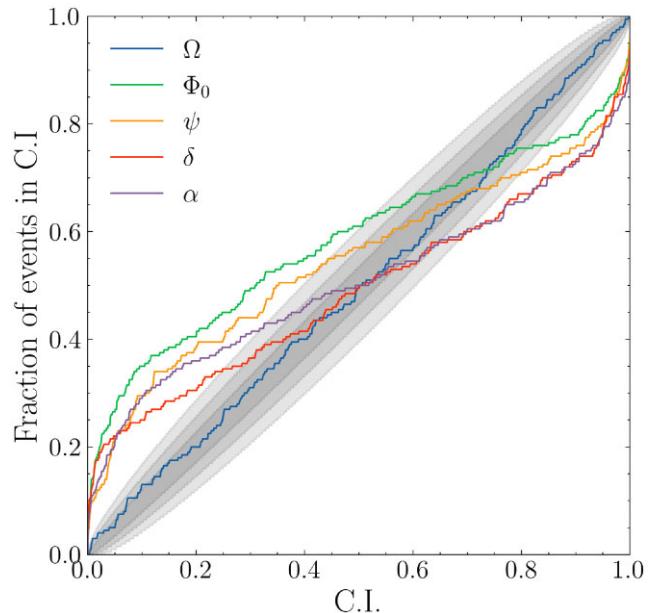


Figure 7. Fraction of injections included within a given credible interval of the estimated posterior, as a function of the credible interval itself (i.e. PP plot). The injections are 200 simulated GW sources generated by drawing randomly five parameters in θ_{gw} from the prior distributions in Table 1. Each coloured curve corresponds to a different parameter (see the legend). The parameters h_0 and ι are fixed at 5×10^{-15} and 1.0 rad, respectively, in order to maintain an approximately constant SNR. The grey shaded contours label the 1σ , 2σ , and 3σ confidence intervals. For parameters with well-estimated posteriors, the PP curve should fall along the diagonal of unit slope. Ω is generally well estimated (i.e. it lies close to the unit diagonal), but the four other parameters show evidence of being overconstrained (i.e. the curves lie above the unit diagonal for low credible intervals, and below the unit diagonal for high credible intervals). This is due to a modelling bias whose origin is discussed in Sections 5.3 and 7.

filter measurement equation as described in Section 5.3 and are well known in the literature (Zhu et al. 2016; Chen & Wang 2022). In this section, we further investigate these biases. In order to elucidate the matter without confusion from the measurement noise, we switch to the high-SNR regime and set $h_0 = 10^{-12}$ (cf. Fig. 6) for the representative system in Table 1.

Fig. 8 displays a corner plot for θ_{gw} , analogous to Fig. 5, but for high SNR ($h_0 = 10^{-12}$). Except for Ω , the injected values lie outside the 90 per cent credible interval, and indeed fall outside the plotted domain. The deviation is most severe for ι with a bias of ≈ 0.3 rad, but is also present to a lesser extent in the static parameters other than Ω . In this section, we demonstrate how the bias stems from dropping the pulsar terms described in Section 5.3.

Fig. 9(a) displays the ι - h_0 contours of the log-likelihood returned by the Kalman filter, that is, $\log \mathcal{L}(\iota, h_0)$, viz. equation (30). The Kalman filter includes the Earth term in the measurement equation, that is, using equation (42). The function $\log \mathcal{L}(\iota, h_0)$ is evaluated across the prior domain and all other parameters are held constant at the true injected value of the representative system in Table 1. The function $\log \mathcal{L}(\iota, h_0)$ is calculated for a single realization of the data and the values are normalized with respect to $\max |\log \mathcal{L}(\iota, h_0)|$. The red point marks the location of $\max [\log \mathcal{L}(\iota, h_0)]$. The orange star marks the location of $\log \mathcal{L}(\iota = 1.0, h_0 = 5 \times 10^{-15})$, that is, the location of the injected parameters (cf. Table 1). The dashed white lines illustrate the (ι, h_0) coordinates of the red point and the orange star. Fig. 9(b) is identical to Fig. 9(a), except that the Kalman filter

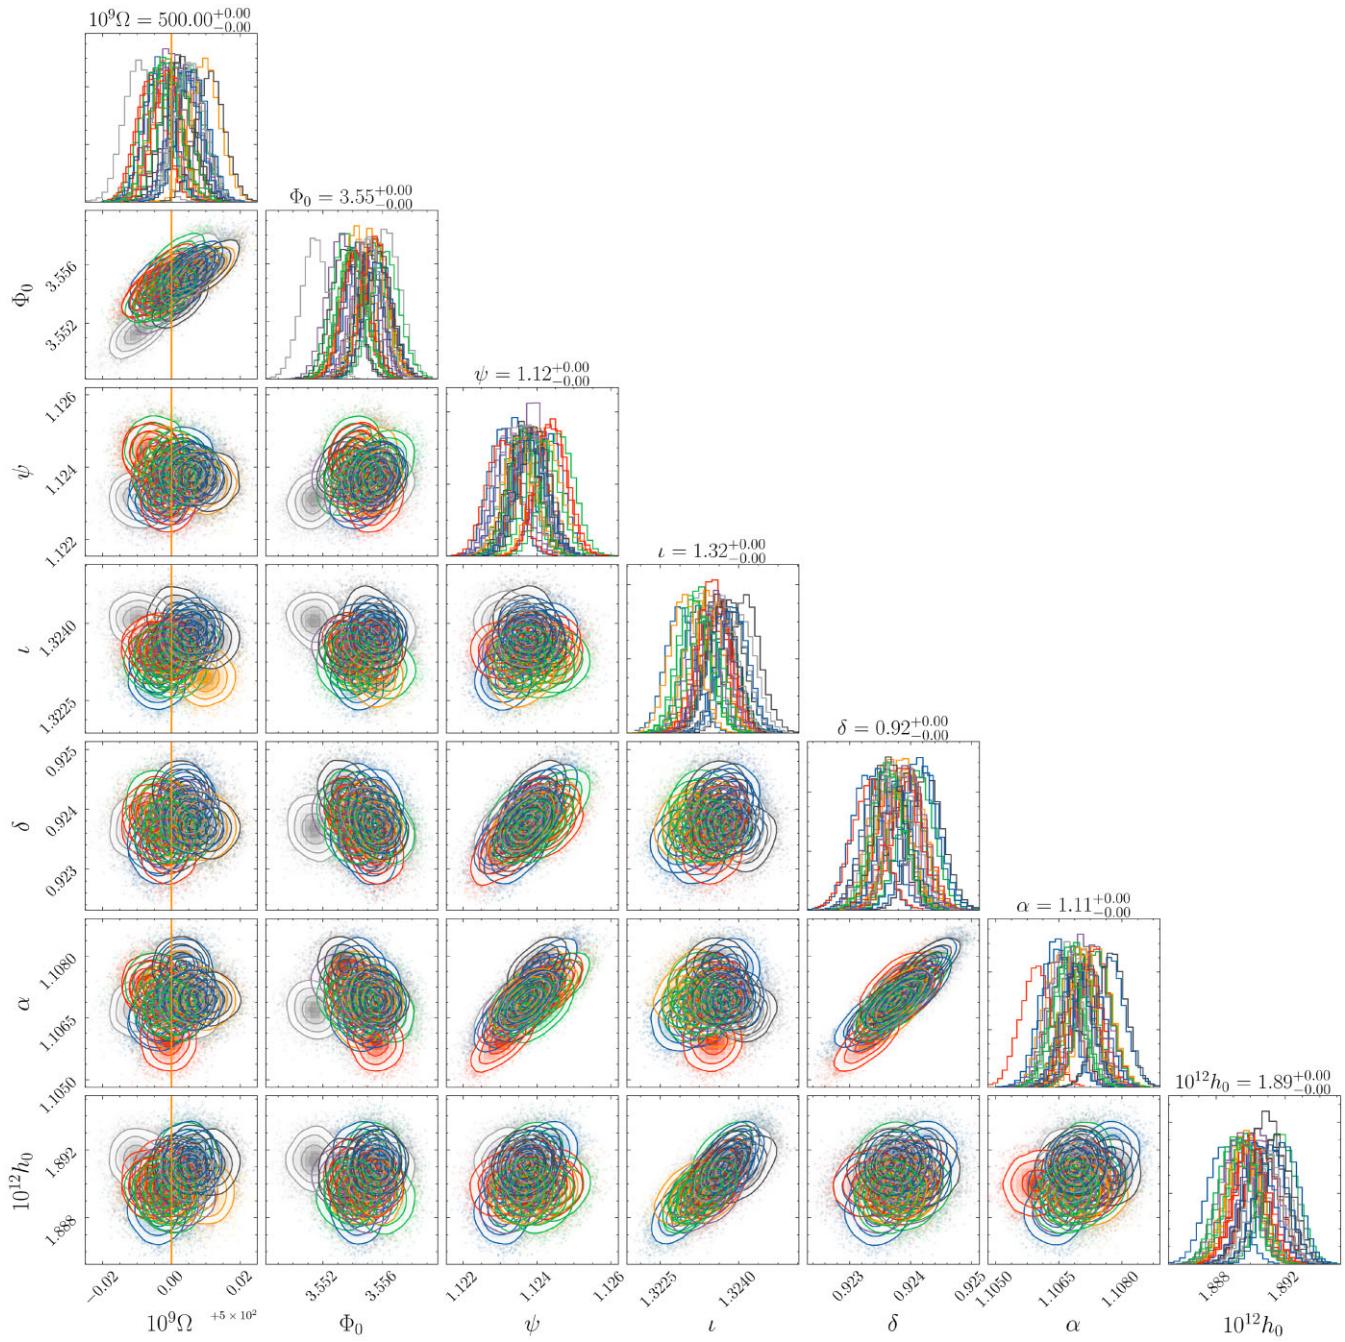


Figure 8. Same as Fig. 5, but for a high-SNR system with $h_0 = 10^{-12}$. For six out of seven parameters the true, injected value does not fall within the 90 per cent credible interval and falls outside the plotted domain. The estimated posteriors are biased away from the injected value due to dropping the pulsar terms. Note that the plotted domain is narrower than in Fig. 5.

includes the Earth and pulsar terms in the measurement equation, that is, using equation (16).

The key observation from Figs 9(a) and (b) is that the location of the maximum ($\iota \approx 1.3$ rad, $\log_{10} h_0 \approx -11.8$) does not exactly coincide with the true, injected value ($\iota = 1.0$ rad, $\log_{10} h_0 = -12$), when only the Earth terms are included. That is, the red point and the orange star occupy different locations in Fig. 9. This is the cause of the bias in the one-dimensional marginalized posteriors in Fig. 8. In contrast, the likelihood maximum coincides with the true, injected value, when the pulsar terms are included, that is, the red point and the orange star are overlaid in Fig. 9(b). Similar results where the

likelihood maximum is offset from the injected values are obtained for the other parameters, with the exception of Ω . An additional example for likelihood contours in the Φ_0 – ψ plane is presented in Fig. 10. The conclusions are analogous to those drawn in the ι – h_0 discussion.

Figs 9(a) and (b) emphasize that the bias observed in Figs 4, 5, and 8 is not a numerical convergence problem, whereby the nested sampler gets stuck in a local optimum. It is instead a structural problem rooted in the fact that the model used to generate the synthetic data is different from the model used for parameter inference. Dropping the pulsar terms biases several of the inferred

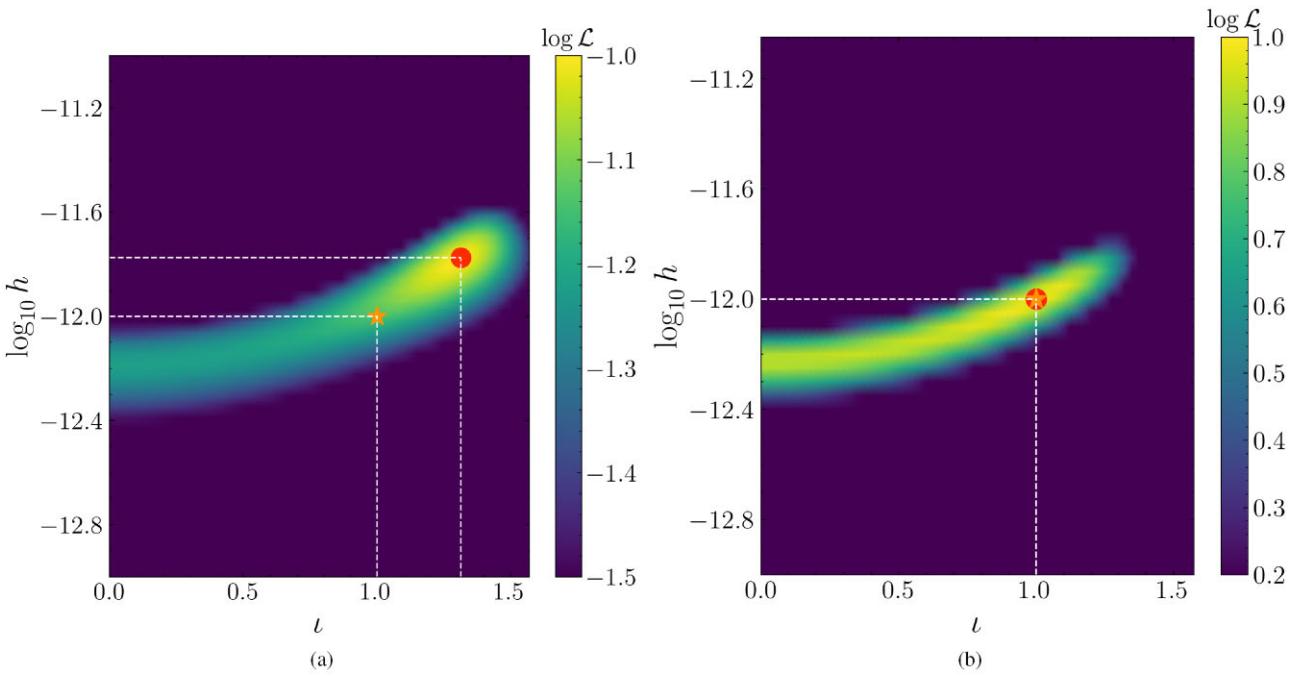


Figure 9. Log-likelihood $\log \mathcal{L}$ (equation 30) contours in the $t-h_0$ plane calculated based on the Kalman filter using (a) exclusively the Earth terms in equation (42) and (b) the Earth terms and the pulsar terms in equation (16). The Kalman filter runs on a single realization of the data. The red point marks the maximum of $\log \mathcal{L}$. The orange star marks the injected parameters, cf. Table 1. The dashed white lines label the (t, h_0) coordinates of the red point and the orange star. In (a), there is a bias; the red point and the orange star occupy different locations. In (b), there is no bias; the red point and the orange star coincide.

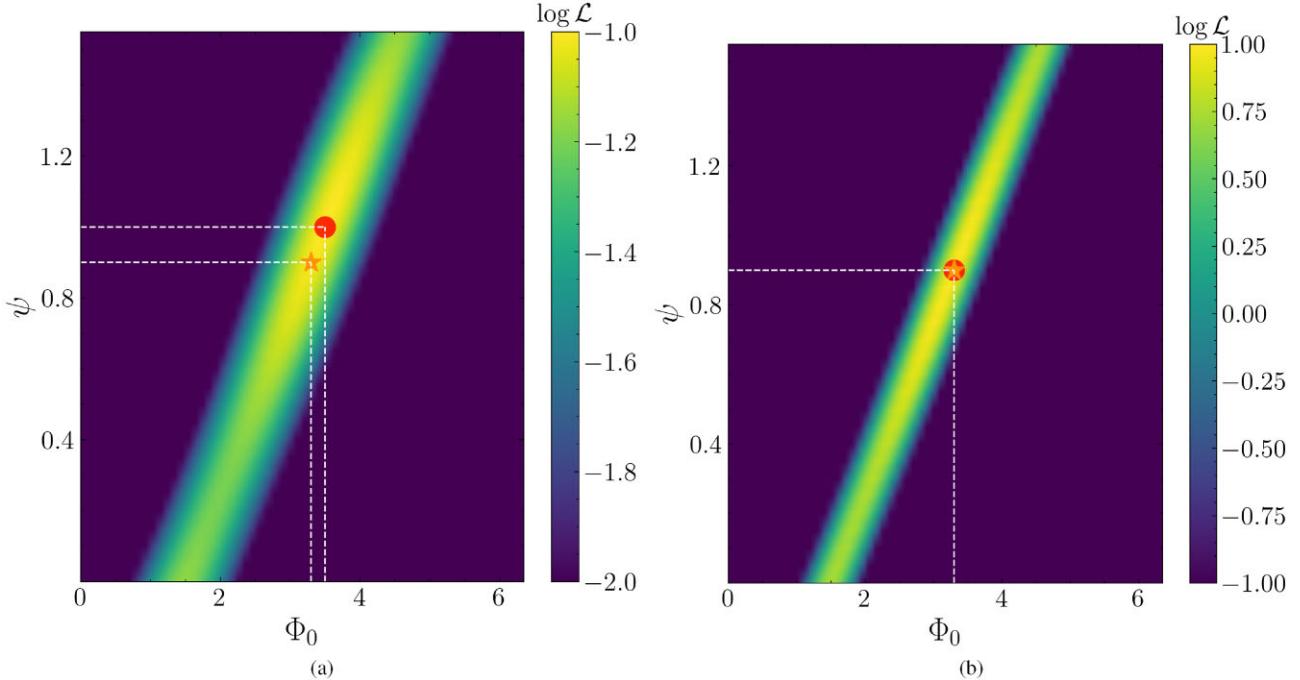


Figure 10. Same as Fig. 9, but for $\log \mathcal{L}$ contours in the $\Phi_0-\psi$ plane. Biases are observed in (a), when the Kalman filter runs using exclusively the Earth terms in equation (42). Including the pulsar term in (b) remedies the bias and the likelihood maximum coincides with the injected static parameter values.

parameters (e.g. those that affect the GW amplitude), not just the sky position (Zhu et al. 2016; Chen & Wang 2022). The bias is not particular to our method; it is shared by all likelihood-based methods that do not include the pulsar terms in the inference model, as is well known in the literature. Quantifying the bias for practical,

astronomical PTA continuous wave searches requires a thorough exploration of the SMBHB parameter space, which falls outside the scope of this introductory paper. The bias also depends on the PTA configuration. If the bias is low enough the uncertainty in the one-dimensional marginalized posterior dominates the bias for quiet

continuous wave sources with low SNR. This is what we observe for the synthetic data; for example, in Fig. 5, the marginalized posterior broadens, as the SNR decreases, until it overwhelms the shift due to the bias. Conversely the bias dominates if it is high enough, or if the source is sufficiently loud. Generalizing the Kalman tracker and nested sampler to correct for the bias will be a key goal of a forthcoming paper.

8 COMPUTATIONAL COST

Bayesian techniques like nested sampling or Markov Chain Monte Carlo are computationally intensive, requiring a large number of likelihood evaluations. In order to readily use these methods, it is important that the central processing unit (CPU) time for a single likelihood evaluation is fast (although some methods for Bayesian inference with expensive likelihoods do exist, e.g. Bilionis & Zabaras 2014; Dinkel et al. 2023; El Gammal et al. 2023). In traditional PTA analyses, the likelihood function is evaluated on all data simultaneously, with corresponding memory demands of large-matrix multiplication and inversion (see e.g. section 7 of Taylor 2021). In contrast, state-space algorithms are iterative and read only the data at the given time-step, with correspondingly smaller memory demands. If state-space algorithms for PTA analysis are to be used in conjunction with Bayesian techniques, it is important that they run quickly. Ideally, they should run at least as fast as traditional PTA analyses, so that they can run in tandem as a cross-check. Moreover, the algorithms must scale favourably with increasing data volume, that is, the number of pulsars in the array, or the number of time samples. In this section, we benchmark the software implementation of the state-space PTA analysis scheme in this paper, as a rough practical guide. We also discuss briefly the scaling of the computational cost as a function of N , the number of pulsars in the PTA, and N_t , the number of TOAs. More extensive benchmarking is postponed, until the analysis scheme is generalized to ingest TOAs instead of $f_m^{(n)}(t)$, and a public, production version of the software is written.

In this paper, as a first pass, we implement a naive version of the Kalman filter, without optimization. As a benchmark, a single likelihood evaluation implemented in PYTHON for the synthetic data presented in Section 4 takes $T_{\mathcal{L}} \sim 9$ ms of CPU time on a 2.6 GHz Intel Core i7 processor. A translation into more performant languages such as C++ (Andrist, Sehr & Garney 2020) or JULIA (Bezanson et al. 2012), the use of PYTHON pre-compilation libraries such as NUMBA (Lam, Pitrou & Seibert 2015) or JAX (Bradbury et al. 2018), or additional optimization (Gorelick & Ozsvárd 2014), would reduce $T_{\mathcal{L}}$.

It is important to understand how $T_{\mathcal{L}}$ scales with both N_t and N . Regarding the former, the theoretical time complexity of the Kalman filter (i.e. its asymptotic behaviour), is $\mathcal{O}(N_t)$, because the Kalman filter is an iterative algorithm. Regarding the latter, the rate-limiting step is set by matrix multiplication, for example, in equations (A4) and (A6) in the Kalman filter's predict and update steps, respectively; see Appendix A1. The dimension of the Kalman filter matrices, for example, equations (A1)–(A9), is set by N . Matrix multiplication without optimization scales as $\mathcal{O}(N^3)$ (Daum 2021). Modern routines for matrix multiplication reduce the complexity to $\sim \mathcal{O}(N^{2.3})$ (Trefethen & Bau 1997). Advanced techniques to further reduce the complexity include CKMS recursion (Kailath, Sayed & Hassibi 2000) or low-rank perturbation methods (Pnevmatikakis et al. 2014). We refer the reader to Raitoharju & Piché (2019) for further information about complexity reduction for the Kalman filter. The

memory complexity of the Kalman filter is $\mathcal{O}(N^2)$, independent of N_t , because the algorithm is iterative.

9 CONCLUSION

In this paper, we demonstrate a new method for the detection and parameter estimation of GWs from individual, monochromatic SMBHBs in PTA data. The new method is complementary to traditional approaches. We track the evolution of the intrinsic pulsar timing noise explicitly via a state-space method, rather than fitting for the ensemble-averaged power spectral density of the noise. That is, we disentangle statistically the specific time-ordered realization of the timing noise from the GW-induced modulations, and thereby infer the GW source parameters conditional on the specific observed realization of the noisy data. We implement a Kalman filter in order to track the intrinsic rotational state evolution of the pulsar and combine it with a Bayesian sampling framework to estimate the posterior distributions of each static parameter, as well as the associated Bayesian evidence (marginalized likelihood) of the model with and without a GW. The favourable time-asymptotic behaviour of the adaptive gain improves the agility of the Kalman tracker compared to alternatives, such as least-squares estimators, and the recursive implementation of the Kalman tracker accelerates the computation.

We test the new method on synthetic data and find that it detects injected signals successfully and estimates their static parameters accurately with relatively low computational cost. We initially focus on a single, astrophysically representative, SMBHB GW source observed synthetically by the 12.5-yr NANOGrav pulsars with $T_{\text{obs}} = 10$ yr. The minimum detectable strain is estimated to be $\min(h_0) \approx 2 \times 10^{-15}$ for $\iota = 1.0$ rad. We then repeat the parameter estimation exercise for 1000 noise realizations to compute the natural dispersion in the recovered values and hence quantify the statistical accuracy of the method in a real astrophysical application, when the true parameter values are unknown. Consistent posteriors are obtained for most realizations. The median WD is limited to $\lesssim 4$ per cent of the width of the prior domain (i.e. negligible railing). The median WD divided by the injected value, an approximate measure of the natural dispersion, ranges from 0.2 per cent for Ω to 25 per cent for h_0 .

Exploration of a broader SMBHB parameter domain at fixed ι and h_0 via 200 randomly sampled parameter vectors reveals a bias in the estimates of each element of the static parameters θ_{gw} , with the exception of Ω . The bias is examined for a specific SMBHB GW source in the limit of high SNR. It is greatest for ι , amounting to ≈ 0.3 rad. Smaller biases of $\lesssim 0.1$ rad are also observed in Φ_0 , ψ , δ , and α . The bias is shown to result from dropping the pulsar terms from the measurement equation in the Kalman filter, consistent with the work of previous authors (Zhu et al. 2016; Chen & Wang 2022). The computational cost of the method is evaluated; a single likelihood evaluation is found to take ~ 9 ms which compares favourably with traditional PTA analyses. The runtime of the full PTA analysis is found to be $\sim 1.5 \times 10^2$ min for the representative SMBHB GW source.

We emphasize that the Kalman tracker and nested sampler in this paper do not supplant traditional PTA analysis approaches; they complement traditional approaches and are most powerful when used in tandem. Relatedly, it is misleading to ask whether the Kalman tracker and nested sampler are more or less sensitive than traditional approaches for two reasons. First, one must still generalize the Kalman tracker and nested sampler to ingest TOAs directly, as happens traditionally, instead of ingesting a frequency

time-series, as in this paper. Generalizing the algorithm is a subtle task and is postponed to a forthcoming paper. Second, the Kalman tracker and nested sampler are conditional on a noise model that is related but different to the noise model in traditional approaches. The analysis in this paper assumes a specific, time-ordered realization of a mean-reverting Ornstein–Uhlenbeck process satisfying equations (1)–(4), whereas traditional analyses assume a stationary Gaussian process described by an ensemble-averaged, power-law power spectral density, whose amplitude and exponent are adjustable. Hence, the sensitivities cannot be compared directly. Clarifying the similarities and differences between various approaches promises to be a fruitful avenue of future work. It is also a subject of attention in audio-band GW data analysis involving hidden Markov models applied to data from terrestrial long-baseline interferometers (Middleton et al. 2020; Abbott et al. 2022a, b).

The approach in this paper can be extended in at least five ways, enumerated as (i)–(v) below.

(i) A natural first extension is to retain the pulsar terms in equation (16). Alternatively, if the method continues to be used with solely the Earth term, in line with standard practice in some published PTA analyses, it would be desirable to evaluate systematically the incurred biases in the model parameters across an astrophysically representative parameter domain, supplementing the results obtained by other authors (Zhu et al. 2016; Chen & Wang 2022; Kimpson et al. 2024).

(ii) In this paper, we consider one specific configuration of the synthetic PTA, namely the same pulsars that make up the 12.5-yr NANOGrav (Section 4.1). It is interesting to compare the performance of the method using different pulsar configurations, for example, those in the PPTA and EPTA. Adding pulsars to PTAs increases the computational cost so it may be advantageous to select a subset of pulsars by exploiting formal optimization techniques from electrical engineering (Speri et al. 2023), although the Kalman filter and nested sampler are already cheaper computationally than some other methods.

(iii) In practice different PTA pulsars are observed with different cadences at different times. Extending the Kalman filter to non-uniform time sampling is straightforward (Zarchan & Musoff 2000).

(iv) The assumption of a monochromatic source is well-justified astrophysically in various regimes (see Section 2.2.1) and is an appropriate starting point for this introductory paper. Nevertheless, SMBHBs are not strictly monochromatic. It is interesting to extend the state-space framework such that f_{gw} evolves in time. For $\Delta f_{\text{gw}} < 1/T_{\text{obs}}$ (see equation 14) evolution adds noise incoherently to the pulsar terms, whilst for $\Delta f_{\text{gw}} > 1/T_{\text{obs}}$ the pulsar terms induce phase shifts that affect the overall phase coherence (Sesana & Vecchio 2010; Perrodin & Sesana 2018). Careful consideration of the evolution of f_{gw} will be needed when including the pulsar terms in the inference model, as in point (i).

(v) Finally, we assume in this paper that there is only one GW source. However, it may be possible to resolve multiple continuous GW sources concurrently (Babak & Sesana 2012). The Kalman filter extends naturally to multiple sources; one can modify equation (16) easily to accommodate a linear superposition of GWs. Taking the logic further, the stochastic background itself is arguably an incoherent sum of many individual GW sources. As long as a way can be found to summarize economically the many static parameters associated with the background sources, it should be possible for a Kalman filter and nested sampler to operate together to detect the stochastic background by generalizing the model selection procedure in Sections 3.3 and 5.5. Summarizing the parameter set economically, while respecting the mathematical structure of the Kalman

filter, is a subtle challenge, which we postpone to a forthcoming paper (Kimpson et al. 2024). If successful, it will complement the traditional approach of cross-correlating pulsar residuals to uncover the Hellings–Downs curve (Hellings & Downs 1983; Agazie et al. 2023a).

ACKNOWLEDGEMENTS

This research was supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), grant no. CE170100004. The numerical calculations were performed on the OzSTAR supercomputer facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government.

DATA AVAILABILITY

No new data were generated or analysed in support of this research.

REFERENCES

- Abbott R. et al., 2022a, *Phys. Rev. D*, 105, 022002
- Abbott R. et al., 2022b, *Phys. Rev. D*, 106, 062002
- Agazie G. et al., 2023a, *ApJ*, 951, L8
- Agazie G. et al., 2023b, *ApJ*, 951, L9
- Agazie G. et al., 2023c, *ApJ*, 951, L50
- Aggarwal K. et al., 2019, *ApJ*, 880, 116
- Alam M. F. et al., 2021, *ApJS*, 252, 5
- Allen B., 1997, in Marck J.-A., Lasota J.-P. (eds), *Relativistic Gravitation and Gravitational Radiation*. Cambridge University Press, Cambridge, UK, p. 373
- Andrist B., Sehr V., Garney B., 2020, *C++ High Performance: Master the Art of Optimizing the Functioning of your C++ Code*. Packt Publishing, <https://books.google.com.au/books?id=tJIREAAQBAJ>
- Antonelli M., Basu A., Haskell B., 2023, *MNRAS*, 520, 2813
- Antoniadis J. et al., 2022, *MNRAS*, 510, 4873
- Antoniadis J. et al., 2023a, *A&A*, 678, A50
- Antoniadis J. et al., 2023b, *A&A*, 678, A48
- Antoniadis J. et al., 2023c, preprint ([arXiv:2306.16226](https://arxiv.org/abs/2306.16226))
- Arjovsky M., Chintala S., Bottou L., 2017, preprint ([arXiv:1701.07875](https://arxiv.org/abs/1701.07875))
- Arzoumanian Z. et al., 2018a, *ApJS*, 235, 37
- Arzoumanian Z. et al., 2018b, *ApJ*, 859, 47
- Arzoumanian Z. et al., 2020a, *ApJ*, 900, 102
- Arzoumanian Z. et al., 2020b, *ApJ*, 905, L34
- Arzoumanian Z. et al., 2023, *ApJ*, 951, L28
- Ashton G. et al., 2022, *Nat. Rev. Methods Primers*, 2, 39
- Ashton G., Talbot C., 2021, *MNRAS*, 507, 2037
- Babak S. et al., 2016, *MNRAS*, 455, 1665
- Babak S., Sesana A., 2012, *Phys. Rev. D*, 85, 044034
- Bailes M. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e028
- Baym G., Pethick C., Pines D., Ruderman M., 1969, *Nature*, 224, 872
- Bécsy B., Cornish N. J., Digman M. C., 2022, *Phys. Rev. D*, 105, 122003
- Bellman R., Åström K., 1970, *Math. Biosci.*, 7, 329
- Bezanson J., Karpinski S., Shah V. B., Edelman A., 2012, preprint ([arXiv:1209.5145](https://arxiv.org/abs/1209.5145))
- Bhagwat S., De Luca V., Franciolini G., Pani P., Riotto A., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 037
- Bilionis I., Zabaras N., 2014, *Inverse Probl.*, 30, 015004
- Bradbury J. et al., 2018, JAX: composable transformations of Python + NumPy programs, <http://github.com/google/jax>
- Brook P. R., Karastergiou A., Buchner S., Roberts S. J., Keith M. J., Johnston S., Shannon R. M., 2013, *ApJ*, 780, L31
- Buchner J., 2021, *J. Open Source Softw.*, 6, 3001
- Buchner J., 2023, *Stat. Surv.*, 17, 169
- Burke-Spolaor S. et al., 2019, *A&AR*, 27, 5

Byrne C., 2005, *Signal Processing: A Mathematical Approach*. Ak Peters Series. Taylor and Francis, Wellesley, MA, <https://books.google.com.au/books?id=IFFyGSowVEkC>

Caballero R. N. et al., 2016, *MNRAS*, 457, 4421

Casella G., Berger R., 2021, *Statistical Inference*. Cengage Learning, <https://books.google.com.au/books?id=FAUVEAAAQBAJ>

Charisi M., Taylor S. R., Witt C. A., Runnoe J., 2024, *Phys. Rev. Lett.*, 132, 061401

Chen J.-W., Wang Y., 2022, *ApJ*, 929, 168

Chen S., Middleton H., Sesana A., Del Pozzo W., Vecchio A., 2017a, *MNRAS*, 468, 404

Chen S., Sesana A., Conselice C. J., 2019, *MNRAS*, 488, 401

Chen S., Sesana A., Del Pozzo W., 2017b, *MNRAS*, 470, 1738

Christensen N., 2019, *Rep. Prog. Phys.*, 82, 016903

Cook S. R., Gelman A., Rubin D. B., 2006, *J. Comput. Graph. Stat.*, 15, 675

Corder G., Foreman D., 2014, *Nonparametric Statistics: A Step-by-Step Approach*. Wiley, New York, <https://books.google.com.au/books?id=hYVYAwAAQBAJ>

Cordes J. M., Greenstein G., 1981, *ApJ*, 245, 1060

Cordes J. M., Lazio T. J. W., 2002, preprint(astro-ph/0207156)

Cordes J. M., Shannon R. M., 2010, preprint (<arXiv:1010.3785>)

D'Alessandro F., McCulloch P. M., Hamilton P. A., Deshpande A. A., 1995, *MNRAS*, 277, 1033

Daum F. E., 2021, *Kalman Filters*. Springer International Publishing, Cham, p. 1067, https://doi.org/10.1007/978-3-030-44184-5_61

Desvignes G. et al., 2016, *MNRAS*, 458, 3341

Dinkel M., Geitner C. M., Robalo Rei G., Nitzler J., Wall W. A., 2024, *Inverse Probl.*, 40, 095008

Dobre S., Bastogne T., Profeta C., Barberi-Heyob M., Richard A., 2012, *Automatica*, 48, 2740, <https://doi.org/10.1016/j.automatica.2012.05.004>

Dudley R. M., 1969, *Ann. Math. Stat.*, 40, 40

Dunn L. et al., 2022, *MNRAS*, 512, 1469

Dunn L., Melatos A., Espinoza C. M., Antonopoulou D., Dodson R., 2023, *MNRAS*, 522, 5469

El Gammal J., Schöneberg N., Torrado J., Fidler C., 2023, *J. Cosmol. Astropart. Phys.*, 2023, 021

Ellis J. A., 2013, *Class. Quantum Gravity*, 30, 224004

Ellis J. A., Cornish N. J., 2016, *Phys. Rev. D*, 93, 084048

Ellis J. A., Siemens X., Creighton J. D. E., 2012, *ApJ*, 756, 175

Espinosa C. M., Antonopoulou D., Dodson R., Stepanova M., Scherer A., 2021, *A&A*, 647, A25

Falxa M. et al., 2023, *MNRAS*, 521, 5077

Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449

Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601

Gardiner C., 2009, *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer Series in Synergetics. Springer, Berlin Heidelberg, <https://books.google.com.au/books?id=otg3PQAAQAAJ>

Gelb A., Kasper J. F., Nash R. A., Price C. F., Sutherland A. A. eds, 1974, *Applied Optimal Estimation*. MIT Press, Cambridge, MA

Gelman A., Carlin J., Stern H., Dunson D., Vehtari A., Rubin D., 2013, *Bayesian Data Analysis*, 3rd edn. Chapman and Hall/CRC Texts in Statistical Science. Taylor and Francis, New York, <https://books.google.com.au/books?id=ZXL6AQAAQBAJ>

Goldreich P., Julian W. H., 1969, *ApJ*, 157, 869

Goldstein J. M., Veitch J., Sesana A., Vecchio A., 2018, *MNRAS*, 477, 5447

Goncharov B. et al., 2021, *MNRAS*, 502, 478

González-Delgado J., Sagar A., Zanon C., Lindorff-Larsen K., Bernadó P., Neuvial P., Cortés J., 2023, *J. Mol. Biol.*, 435, 168053 <https://doi.org/10.1016/j.jmb.2023.168053>

Gorelick M., Ozsvárd I., 2014, *High Performance Python: Practical Performant Programming for Humans*. O'Reilly Media, Incorporated, Sebastopol, CA, <https://books.google.com.au/books?id=QSsOBQAAQBAJ>

Greenstein G., 1970, *Nature*, 227, 791

Gügercioğlu E., Alpar M. A., 2017, *MNRAS*, 471, 4827

Guillaume J. H. et al., 2019, *Environ. Model. Softw.*, 119, 418, <https://doi.org/10.1016/j.envsoft.2019.07.007>

Handley W. J., Hobson M. P., Lasenby A. N., 2015, *MNRAS*, 450, L61

Hellings R. W., Downs G. S., 1983, *ApJ*, 265, L39

Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655

Hobbs G., Dai S., Manchester R. N., Shannon R. M., Kerr M., Lee K.-J., Xu R.-X., 2019, *Res. Astron. Astrophys.*, 19, 020

Hobbs G., Lyne A. G., Kramer M., 2010, *MNRAS*, 402, 1027

Jaffe A. H., Backer D. C., 2003, *ApJ*, 583, 616

Jankowski F. et al., 2018, *MNRAS*, 484, 3691

Janet F. A., Lommen A., Larson S. L., Wen L., 2004, *ApJ*, 606, 799

Johnson A. D. et al., 2024, *Phys. Rev. D*, 109, 103012

Johnston S., Galloway D., 1999, *MNRAS*, 306, L50

Jones P. B., 1990, *MNRAS*, 246, 364

Kailath T., Sayed A., Hassibi B., 2000, *Linear Estimation*. Prentice-Hall Information and System Sciences Series. Prentice Hall, <https://books.google.com.au/books?id=zNJFAQAAIAJ>

Kalman R. E., 1960, *J. Basic Eng.*, 82, 35

Kerr M., Hobbs G., Johnston S., Shannon R. M., 2015, *MNRAS*, 455, 1845

Kimpson T. et al., 2024, *MNRAS*, submitted

Kimpson T., Paxton E. A., Chantry M., Palmer T., 2023, *Q. J. R. Meteorol. Soc.*, 149, 843

Kramer M., Lyne A. G., O'Brien J. T., Jordan C. A., Lorimer D. R., 2006, *Science*, 312, 549

Kuhn H. W., 1955, *Nav. Res. Logisti. Q.*, 2, 83, <https://doi.org/10.1002/nav.3800020109>

Lam M. T. et al., 2019, *ApJ*, 872, 193

Lam S. K., Pitrou A., Seibert S., 2015, in *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC*. ACM, p. 1

Lasky P. D., Melatos A., Ravi V., Hobbs G., 2015, *MNRAS*, 449, 3293

Lee K. J., Wex N., Kramer M., Stappers B. W., Bassa C. G., Janssen G. H., Karuppusamy R., Smits R., 2011, *MNRAS*, 414, 3251

Lentati L. et al., 2015, *MNRAS*, 453, 2576

Lentati L., Alexander P., Hobson M. P., Taylor S., Gair J., Balan S. T., van Haasteren R., 2013, *Phys. Rev. D*, 87, 104021

Lewis F., Xie L., Popa D., 2017, *Optimal and Robust Estimation: With an Introduction to Stochastic Control Theory*, Second Edition. CRC Press, Boca Raton, <https://books.google.com.au/books?id=hNXLBQAAQBAJ>

Lower M. E. et al., 2021, *MNRAS*, 508, 3251

Luo J. et al., 2021, *ApJ*, 911, 45

Lyne A., Hobbs G., Kramer M., Stairs I., Stappers B., 2010, *Science*, 329, 408

Maggiore M., 2018, *Gravitational Waves: Volume 2: Astrophysics and Cosmology*. Oxford University Press, Oxford, <https://doi.org/10.1093/oso/9780198570899.001.0001>

Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, *AJ*, 129, 1993

McWilliams S. T., Ostriker J. P., Pretorius F., 2014, *ApJ*, 789, 156

Melatos A., Dunn L. M., Suvorova S., Moran W., Evans R. J., 2020, *ApJ*, 896, 78

Melatos A., Link B., 2014, *MNRAS*, 437, 21

Melatos A., O'Neill N. J., Meyers P. M., O'Leary J., 2023, *ApJ*, 944, 64

Melatos A., Peralta C., Wyithe J. S. B., 2008, *ApJ*, 672, 1103

Meyers P. M., Melatos A., O'Neill N. J., 2021a, *MNRAS*, 502, 3113

Meyers P. M., O'Neill N. J., Melatos A., Evans R. J., 2021b, *MNRAS*, 506, 3349

Middleton H., Clearwater P., Melatos A., Dunn L., 2020, *Phys. Rev. D*, 102, 023006

Morris M. J., Lipp A. G., Roberts G. G., 2023, *Geophys. Res. Lett.*, 50, e2023GL103880

Mukherjee P., Parkinson D., Liddle A. R., 2006, *ApJ*, 638, L51

O'Neill N. J., Meyers P. M., Melatos A., 2024, *MNRAS*, 530, 4648

Pártay L. B., Bartók A. P., Csányi G., 2009, preprint (<arXiv:0906.3544>)

Parthasarathy A. et al., 2021, *MNRAS*, 502, 407

Paxton E. A., Chantry M., Klöwer M., Saffin L., Palmer T., 2022, *J. Clim.*, 35, 1215

Pennucci T. T., Collaboration T. N., 2020, *The NANOGrav 12.5-year Wideband Data Set (version 12yv4)*, Zenodo, <https://doi.org/10.5281/zenodo.4312887>

Peralta C., Melatos A., Giacobello M., Ooi A., 2006, *ApJ*, 651, 1079

Perera B. B. P. et al., 2019, *MNRAS*, 490, 4666

Perrodin D., Sesana A., 2018, in Rezzolla L., Pizzochero P., Jones D. I., Rea N., Vidaña Leds, *Astrophysics and Space Science Library*, Vol. 457, *The Physics and Astrophysics of Neutron Stars*. Springer Nature Switzerland AG, Switzerland, p. 95

Petiteau A., Babak S., Sesana A., de Araújo M., 2013, *Phys. Rev. D*, 87, 064036

Pitkin M., 2018, *J. Open Source Softw.*, 3, 538

Pnevmatikakis E. A., Rad K. R., Huggins J., Paninski L., 2014, *J. Comput. Graph. Stat.*, 23, 316

Price S., Link B., Shore S. N., Nice D. J., 2012, *MNRAS*, 426, 2507

Prix R., Krishnan B., 2009, *Class. Quantum Gravity*, 26, 204013

Raitoharju M., Piché R., 2019, *IEEE Aerosp. Electron. Syst. Mag.*, 34, 2

Rajagopal M., Romani R. W., 1995, *ApJ*, 446, 543

Ravi V., Wyithe J. S. B., Hobbs G., Shannon R. M., Manchester R. N., Yardley D. R. B., Keith M. J., 2012, *ApJ*, 761, 84

Ravi V., Wyithe J. S. B., Shannon R. M., Hobbs G., 2015, *MNRAS*, 447, 2772

Reardon D. J. et al., 2023, *ApJ*, 951, L6

Renzini A. I., Goncharov B., Jenkins A. C., Meyers P. M., 2022, *Galaxies*, 10, 34

Rosado P. A., Sesana A., Gair J., 2015, *MNRAS*, 451, 2417

Sampson L., Cornish N. J., McWilliams S. T., 2015, *Phys. Rev. D*, 91, 084055

Sanidas S. A., Battye R. A., Stappers B. W., 2012, *Phys. Rev. D*, 85, 122003

Särkkä S., 2013, *Bayesian Filtering and Smoothing*. Cambridge University Press, Cambridge, UK, <https://books.google.com.au/books?id=5VlsAQAQBAJ>

Schutz B., 2022, *Fundamentals of Gravitational Radiation*, 3rd edn. Cambridge University Press, Cambridge, UK, p. 202

Sesana A., 2013, *Class. Quantum Gravity*, 30, 224014

Sesana A., Vecchio A., 2010, *Phys. Rev. D*, 81, 104008

Sesana A., Vecchio A., Colacino C. N., 2008, *MNRAS*, 390, 192

Shannon R. M. et al., 2013, *ApJ*, 766, 5

Shannon R. M., Cordes J. M., 2010, *ApJ*, 725, 1607

Simon D., 2006, *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. Wiley-Interscience, New York, NY

Skilling J., 2006, *Bayesian Anal.*, 1, 833

Sornette D., 2004, *Critical Phenomena in Natural Sciences: Chaos, Fractals Self-organization and Disorder: Concepts and Tools*. Springer Series in Synergetics. Springer, Heidelberg

Speagle J. S., 2020, *MNRAS*, 493, 3132

Speri L., Porayko N. K., Falxa M., Chen S., Gair J. R., Sesana A., Taylor S. R., 2023, *MNRAS*, 518, 1802

Spiewak R. et al., 2022, *Publ. Astron. Soc. Aust.*, 39, e027

Stairs I. H. et al., 2019, *MNRAS*, 485, 3230

Stairs I. H., Lyne A. G., Shemar S. L., 2000, *Nature*, 406, 484

Sykes B., Middleton H., Melatos A., Di Matteo T., DeGraf C., Bhowmick A., 2022, *MNRAS*, 511, 5241

Tarafdar P. et al., 2022, *Publ. Astron. Soc. Aust.*, 39, e053

Taylor S. R., 2021, preprint ([arXiv:2105.13270](https://arxiv.org/abs/2105.13270))

Taylor S. R., Huerta E. A., Gair J. R., McWilliams S. T., 2016, *ApJ*, 817, 70

Taylor S. R., Simon J., Sampson L., 2017, *Phys. Rev. Lett.*, 118, 181102

Tiburzi C., 2018, *Publ. Astron. Soc. Aust.*, 35, e013

Trassinelli M., 2019, *Proceedings*, 33, 14

Trefethen L., Bau D., 1997, *Numerical Linear Algebra*. Other Titles in Applied Mathematics, Society for Industrial and Applied Mathematics, https://books.google.com.au/books?id=4Mou5YpRD_kC

Urama J. O., Link B., Weisberg J. M., 2006, *MNRAS*, 370, L76

van Eysden C. A., Melatos A., 2010, *MNRAS*, 409, 1253

Vargas A., Melatos A., 2023, *TBD*, 1, 1

Vaserstein L., 1969, *Probl. Peredachi Inf.*, 5, 64

Verbiest J. P. W., Oslowski S., Burke-Spolaor S., 2021, in Bambi C., Katsanevas S., Kokkotas K.D., eds, *Handbook of Gravitational Wave Astronomy*. Springer, Berlin, p. 4

Verbiest J. P. W., Weisberg J. M., Chael A. A., Lee K. J., Lorimer D. R., 2012, *ApJ*, 755, 39

Villani C., 2003, *Topics in Optimal Transportation*. Graduate Studies in Mathematics, American Mathematical Society, <https://books.google.com.au/books?id=idyFAwAAQBAJ>

Villani C., 2009, *The Wasserstein Distances*. Springer, Berlin, Heidelberg, p. 93, https://doi.org/10.1007/978-3-540-71050-9_6

Walter E., Pronzato L., 1996, *Math. Comput. Simul.*, 42, 125, [https://doi.org/10.1016/0378-4754\(95\)00123-9](https://doi.org/10.1016/0378-4754(95)00123-9)

Wan E., Van Der Merwe R., 2000, in *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium* (Cat. No.00EX373). p. 153

Wyithe J. S. B., Loeb A., 2003, *ApJ*, 590, 691

Xu H. et al., 2023, *Res. Astron. Astrophys.*, 23, 075024

Xue X. et al., 2021, *Phys. Rev. Lett.*, 127, 251303

Yao J. M., Manchester R. N., Wang N., 2017, *ApJ*, 835, 29

Yardley D. R. B. et al., 2010, *MNRAS*, 407, 669

Zarchan P., Musoff H., 2000, *Fundamentals of Kalman Filtering: A Practical Approach*. Progress in Astronautics and Aeronautics. American Institute of Aeronautics and Astronautics, Reston, VA, <https://books.google.com.au/books?id=AQxRAAAAMAAJ>

Zhu X. J. et al., 2014, *MNRAS*, 444, 3709

Zhu X.-J. et al., 2015a, *MNRAS*, 449, 1650

Zhu X.-J. et al., 2015b, *MNRAS*, 449, 1650

Zhu X.-J., Wen L., Xiong J., Xu Y., Wang Y., Mohanty S. D., Hobbs G., Manchester R. N., 2016, *MNRAS*, 461, 1317

Zic A. et al., 2023, *Publ. Astron. Soc. Aust.*, 40, e049

APPENDIX A: KALMAN FILTER

In this appendix, we describe the Kalman filter algorithm used in this paper. General recursion relations for the discrete-time Kalman filter are written down for an arbitrary linear dynamical system in Appendix A1. The mapping onto the specific continuous-time state-space model in Section 2 is written down in Appendix A2. Separately, in Appendix B, we compare the Kalman filter and Ornstein–Uhlenbeck model with traditional PTA data analysis techniques.

A1 Recursion equations

The linear Kalman filter operates on temporally discrete, noisy measurements \mathbf{Y}_k , which are related to a set of unobservable discrete system states \mathbf{X}_k , via a linear transformation

$$\mathbf{Y}_k = \mathbf{H}_k \mathbf{X}_k + \mathbf{v}_k, \quad (A1)$$

where \mathbf{H}_k is the measurement matrix or observation model, \mathbf{v}_k is a zero-mean Gaussian measurement noise, $\mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R}_k)$ with covariance \mathbf{R}_k , and the subscript k labels the time-step. The Kalman filter evolves the underlying states according to

$$\mathbf{X}_k = \mathbf{F}_k \mathbf{X}_{k-1} + \mathbf{G}_k \mathbf{u}_k + \mathbf{w}_k, \quad (A2)$$

where \mathbf{F}_k is the system dynamics matrix, \mathbf{G}_k is the control matrix, \mathbf{u}_k is the control vector, and \mathbf{w}_k is a zero-mean Gaussian process noise, $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q}_k)$ with covariance \mathbf{Q}_k

The Kalman filter is a recursive estimator with two distinct stages: a ‘predict’ stage and an ‘update’ stage. The predict stage predicts $\hat{\mathbf{X}}_{k|k-1}$, the estimate of the state at discrete step k , given the state estimates from step $k-1$. Specifically, the predict step proceeds as

$$\hat{\mathbf{X}}_{k|k-1} = \mathbf{F}_k \hat{\mathbf{X}}_{k-1|k-1} + \mathbf{G}_k \mathbf{u}_k, \quad (A3)$$

$$\hat{\mathbf{P}}_{k|k-1} = \mathbf{F}_k \hat{\mathbf{P}}_{k-1|k-1} \mathbf{F}_k^\top + \mathbf{Q}_k, \quad (A4)$$

where $\hat{\mathbf{P}}_{k|k-1}$ is the covariance of the prediction. Note that the predict stage is independent of the measurements. The measurement \mathbf{Y}_k is included to update the prediction during the update stage as follows:

$$\mathbf{e}_k = \mathbf{Y}_k - \mathbf{H}_k \hat{\mathbf{X}}_{k|k-1}, \quad (A5)$$

$$\mathbf{S}_k = \mathbf{H}_k \hat{\mathbf{P}}_{k|k-1} \mathbf{H}_k^\top + \mathbf{R}_k, \quad (A6)$$

$$\mathbf{K}_k = \hat{\mathbf{P}}_{k|k-1} \mathbf{H}_k^\top \mathbf{S}_k^{-1}, \quad (\text{A7})$$

$$\hat{\mathbf{X}}_{k|k} = \hat{\mathbf{X}}_{k|k-1} + \mathbf{K}_k \boldsymbol{\epsilon}_k, \quad (\text{A8})$$

$$\hat{\mathbf{P}}_{k|k} = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \hat{\mathbf{P}}_{k|k-1}. \quad (\text{A9})$$

Equation (A7) defines the Kalman gain \mathbf{K}_k which is defined so as to minimize the mean-squared error in the state estimate, that is, $\mathbf{K}_k = \text{argmin} \{ \mathbf{E}[(\mathbf{X}_k - \hat{\mathbf{X}}_k)^2] \}$. For full reviews of the Kalman filter, including its derivation, we refer the reader to Gelb et al. (1974) and Zarchan & Musoff (2000).

To apply the Kalman filter in practice means specifying the eight component matrices that make up the ‘machinery’ of the filter: \mathbf{X}_k , \mathbf{Y}_k , \mathbf{F}_k , \mathbf{G}_k , \mathbf{u}_k , \mathbf{H}_k , \mathbf{Q}_k , and \mathbf{R}_k . In Appendix A2, we describe how the machinery is defined for the state-space model in Section 2.

A2 State-space representation of a PTA analysis

We apply the Kalman recursion relations in Appendix A1 to the state-space model of a PTA with N pulsars described in Section 2 as follows.

We identify $\mathbf{X}(t)$ with a vector of length N composed of the intrinsic pulsar frequency states, that is,

$$\mathbf{X}(t) = (f_p^{(1)}(t), f_p^{(2)}(t), \dots, f_p^{(N)}(t)). \quad (\text{A10})$$

Analogously, we package the measured pulsar frequencies as

$$\mathbf{Y}(t) = (f_m^{(1)}(t), f_m^{(2)}(t), \dots, f_m^{(N)}(t)). \quad (\text{A11})$$

The states evolve according to the continuous stochastic differential equation (cf. equation 2)

$$d\mathbf{X} = \mathbf{A}\mathbf{X}dt + \mathbf{C}(t)dt + \boldsymbol{\Sigma}d\mathbf{B}(t), \quad (\text{A12})$$

where \mathbf{A} is a diagonal $N \times N$ matrix,

$$\mathbf{A} = \text{diag}(-\gamma^{(1)}, -\gamma^{(2)}, \dots, -\gamma^{(N)}), \quad (\text{A13})$$

and $\mathbf{C}(t)$ is a time-dependent vector with n th component

$$C^{(n)} = \gamma^{(n)} [f_{\text{em}}^{(n)}(t_1) + \dot{f}_{\text{em}}^{(n)}(t_1)t] + \dot{f}_{\text{em}}^{(n)}(t_1). \quad (\text{A14})$$

The $N \times N$ square matrix $\boldsymbol{\Sigma}$ governs the magnitude of the increments of Brownian motion (Wiener process) $d\mathbf{B}(t)$, with

$$\boldsymbol{\Sigma} = \text{diag}(\sigma^{(1)}, \sigma^{(2)}, \dots, \sigma^{(N)}). \quad (\text{A15})$$

In the idealized model above, each pulsar’s rotational state evolves phenomenologically according to a mean-reverting Ornstein–Uhlenbeck process, described by a Langevin equation (equation A12), whose general solution is given by (Gardiner 2009)

$$\mathbf{X}(t) = e^{\mathbf{A}t} \mathbf{X}(0) + \int_0^t e^{\mathbf{A}(t-t')} \mathbf{C}(t') dt' + \int_0^t e^{\mathbf{A}(t-t')} \boldsymbol{\Sigma} d\mathbf{B}(t'). \quad (\text{A16})$$

From equation (A16), we construct the discrete, recursive solution for $\mathbf{X}(t_k) = \mathbf{X}_k$ in the form of equation (A2), with

$$\mathbf{F}_k = e^{\mathbf{A}\Delta t} \quad (\text{A17})$$

$$= \text{diag}(e^{-\gamma^{(1)}\Delta t}, e^{-\gamma^{(2)}\Delta t}, \dots, e^{-\gamma^{(N)}\Delta t}), \quad (\text{A18})$$

$$\mathbf{G}_k \mathbf{u}_k = \int_{t_k}^{t_{k+1}} e^{\mathbf{A}(t_{k+1}-t')} \mathbf{C}(t') dt', \quad (\text{A19})$$

$$= (G_k^{(1)}, G_k^{(2)}, \dots, G_k^{(N)}), \quad (\text{A20})$$

$$\mathbf{w}_k = \int_{t_k}^{t_{k+1}} e^{\mathbf{A}(t_{k+1}-t')} \boldsymbol{\Sigma} d\mathbf{B}(t'), \quad (\text{A21})$$

$$G_k^{(n)} = f_{\text{em}}^{(n)}(t_1) + \dot{f}_{\text{em}}^{(n)}(t_1)(\Delta t + t_k) - e^{-\gamma^{(n)}\Delta t} [f_{\text{em}}^{(n)}(t_1) + \dot{f}_{\text{em}}^{(n)}(t_1)t_k], \quad (\text{A22})$$

and $\Delta t = t_{k+1} - t_k$. From equation (A21), the process noise covariance matrix is

$$\mathbf{Q}_k \delta_{kj} = \langle \boldsymbol{\eta}_k \boldsymbol{\eta}_j^\top \rangle \quad (\text{A23})$$

$$= \text{diag}(\mathbf{Q}^{(1)}, \mathbf{Q}^{(2)}, \dots, \mathbf{Q}^{(N)}), \quad (\text{A24})$$

with

$$\mathbf{Q}^{(n)} = \frac{[\sigma^{(n)}]^2}{2\gamma^{(n)}} \left[1 - e^{-2\gamma^{(n)}\Delta t} \right]. \quad (\text{A25})$$

The Einstein summation convention is suspended temporarily in the left-hand side of equation (A25). The two remaining unspecified component matrices of the Kalman filter are the measurement matrix \mathbf{H}_k and the measurement covariance matrix \mathbf{R}_k . These are defined straightforwardly from equations (16)–(21). Specifically, \mathbf{H}_k is a diagonal matrix where the n th component of the diagonal is given by $g^{(n)}(t_k)$ from equation (16). The measurement covariance satisfies $\mathbf{R}_k = \mathbf{E}[\mathbf{v}\mathbf{v}^\top] = \sigma_m^2$ for all k .

APPENDIX B: COMPARISON WITH TRADITIONAL PTA ANALYSES

In this appendix, we compare the state-space formulation of PTA data analysis described in this paper with traditional formulations. In Appendix B1, we focus on one key difference: the asymptotic behaviour of the adaptive gain in the Kalman filter compared to traditional, least-squares estimators. The adaptive gain controls the fraction of new information that is incorporated into the updated state estimate at every time-step and hence controls how nimbly the tracking scheme responds to new data. We illustrate the effect of the adaptive gain on a simple Wiener process as a pedagogical example and then explain how it applies analogously to the PTA problem. In Appendix B2, we compare the Ornstein–Uhlenbeck description of the process noise intrinsic to the pulsar, specified by equations (2)–(4), with the traditional approach in terms of modelling red noise as a Gaussian process with a (broken) power-law power spectral density. We show that the two descriptions are equivalent in certain limits.

B1 Adaptive gain

Many traditional PTA analyses fit timing data to a phase model by least-squares estimation. The state-space scheme in this paper achieves the same goal (for a pulse frequency model rather than a phase model, strictly speaking) using a Kalman filter. It is natural to ask how the two approaches differ, if at all.

To understand the difference from first principles, consider as a simplified pedagogical example the problem of estimating an unobservable state X using a measurement Y , where X is generated by a Wiener process, viz.

$$\dot{X} = w, \quad (\text{B1})$$

$$Y = X + v, \quad (\text{B2})$$

with $w \sim \mathcal{N}(0, Q)$ and $v \sim \mathcal{N}(0, R)$, and where an overdot denotes a derivative with respect to time. We now compare the least-squares estimator and the Kalman filter for the simplified system in equations

(B1) and (B2) in two cases: (i) $Q = 0$ (i.e. the state is constant with no process noise) and (ii) $Q \neq 0$. In order to compare directly the Kalman filter (which is recursive) and the least-squares estimator, we reformulate the batch least-squares regression as a recursive least-squares filter. Batch and recursive least squares are equivalent; both methods minimize the sum of the squared errors and produce identical estimates of the state, \hat{X} .

First, consider $Q = 0$. In recursive least-squares, the unbiased minimum variate estimate of the state at time-step k is

$$\hat{X}_{k,RLS} = \frac{1}{k} \sum_{i=1}^k Y_i. \quad (\text{B3})$$

When a new data point is ingested at time-step $k + 1$, the estimate of the state is updated according to

$$\hat{X}_{k+1,RLS} = \frac{1}{k+1} \sum_{i=1}^{k+1} Y_i \quad (\text{B4})$$

$$= \hat{X}_{k,RLS} + \frac{1}{k+1} (Y_{k+1} - \hat{X}_{k,RLS}). \quad (\text{B5})$$

The prefactor $(k+1)^{-1}$ in equation (B5) is the recursive least-squares gain (cf. equation A7) and the quantity $Y_{k+1} - \hat{X}_{k,RLS}$ is the innovation (cf. equation A5). As the number of data points increases the gain tends to zero; in recursive least-squares the additional datum Y_{k+1} becomes less and less important, as k increases and more and more observations are obtained. As the gain tends to zero, the new observations are asymptotically downweighted.

Now apply Kalman filtering to $Q = 0$ instead. The state estimate is updated in response to new data according to

$$\hat{X}_{k+1,KF} = \hat{X}_{k,KF} + \frac{P_0}{R + kP_0} (Y_{k+1} - \hat{X}_{k,KF}), \quad (\text{B6})$$

where P_0 is the initial covariance in \hat{X} . In equation (B6), the quantity $P_0/(R + kP_0)$ is the Kalman gain. As with recursive least-squares, the gain tends to zero as k increases. Hence, in the special case $Q = 0$, the recursive least-squares and Kalman filtering behave equivalently in the long term.

Now consider $Q \neq 0$. The recursive least-squares estimate of the state proceeds as before via equation (B5). What about the Kalman filter? For the system described by equations (B1) and (B2), the continuous-time Riccati equation for the propagation of the error covariance $P(t)$ is (Lewis, Xie & Popa 2017)

$$\dot{P} = Q - \frac{P(t)^2}{R}. \quad (\text{B7})$$

The continuous-time Kalman gain is

$$K(t) = \frac{P(t)}{R}. \quad (\text{B8})$$

The solution to equation (B7) is

$$P(t) = \sqrt{RQ} \left[\frac{P_0 \cosh(t\sqrt{Q/R}) + \sqrt{QR} \sinh(t\sqrt{Q/R})}{P_0 \sinh(t\sqrt{Q/R}) + \sqrt{QR} \cosh(t\sqrt{Q/R})} \right]. \quad (\text{B9})$$

From equations (B8) and (B9), we find $K(t) \rightarrow \sqrt{Q/R} > 0$ as t increases (equivalent to increasing k in the discrete case). That is, the gain tends to a positive definite value, unlike for recursive least squares. Hence, the Kalman filter is more responsive asymptotically to additional data than a least-squares estimator. Informally speaking, it is a more nimble tracker, which explains its preferred status in many electrical and mechanical engineering applications (Gelb et al. 1974; Zarchan & Musoff 2000; Byrne 2005; Särkkä 2013).

PTA data analysis is far more complicated than the pedagogical example described by equations (B1) and (B2). Fundamentally, though, the timing noise tracking step can be understood in the same way, as a state estimation problem. In this paper, state estimation is performed with a Kalman filter. In traditional analyses, it is performed through least-squares estimation using timing software such as TEMPO2 (Hobbs et al. 2006) or PINT (Luo et al. 2021). Specifically, traditional least-squares analyses fit a model of the TOA residuals vector δt of the form

$$\delta t = \mathbf{M}\epsilon + \mathcal{F}\mathbf{a} + \mathbf{n} \quad (\text{B10})$$

where $\mathbf{M}\epsilon$ represents the deterministic deviation from the least-squares fit to a Taylor-series phase model, $\mathcal{F}\mathbf{a}$ represents the stochastic red noise modelled as a sum of sine and cosine Fourier modes with amplitudes drawn from a power spectral density (usually of power-law form), and \mathbf{n} represents the stochastic white noise component, with covariance matrix N (Taylor 2021; Johnson et al. 2024). Minimizing the squared error is equivalent to maximizing the associated likelihood

$$p(\mathbf{r}|\epsilon, \mathbf{a}) = \frac{1}{\sqrt{\det(2\pi N)}} \exp\left(-\frac{1}{2}\mathbf{r}^T N \mathbf{r}\right), \quad (\text{B11})$$

with $\mathbf{r} = \delta t - \mathbf{M}\epsilon - \mathcal{F}\mathbf{a}$. The above formulation is analogous to the $Q \neq 0$ case analysed above, and the conclusion is the same: the least-squares gain tends to zero asymptotically, whereas the Kalman gain does not. This represents a difference between traditional PTA analyses and the approach in this paper. We refer the reader to Taylor (2021) for additional details on traditional PTA data analysis methods.

B2 Red noise power spectral density

In Section 2, the intrinsic achromatic spin wandering of the pulsar, $f_p^{(n)}(t)$, is modelled as an Ornstein–Uhlenbeck process governed by equations (1)–(4). The Ornstein–Uhlenbeck model captures the main qualitative features of a typical PTA pulsar, namely a deterministic secular spin down perturbed by stochastic, small-amplitude, mean-reverting fluctuations.

In traditional PTA analyses, the red-spectrum timing noise fluctuations are described as a zero-mean Gaussian random process and modelled via a finite decomposition into a Fourier basis (see Appendix B1). The Fourier coefficients are determined by the power spectral density of the TOA residuals, which is assumed to take the standard form (e.g. Goncharov et al. 2021; Taylor 2021),

$$\rho(f) = \frac{A_a^2}{12\pi^2} \frac{1}{T_{\text{obs}}} \left(\frac{f}{1\text{yr}^{-1}} \right)^{-\gamma} \text{yr}^2, \quad (\text{B12})$$

where f denotes the Fourier frequency. The power spectral density is defined by two hyperparameters, the amplitude A_a and the exponent γ . Traditional PTA analyses seek to estimate the values of the hyperparameters, rather than the Fourier coefficients themselves, which are drawn randomly by treating equation (B12) as a probability density function.

Equations (1)–(4) are equivalent to specifying a particular form of power spectral density, which is related to but different from equation (B12). Specifically, equations (1)–(4) lead to a broken power law with $\rho(f) \propto f^{-4}$ for $f \gtrsim \gamma^{(n)}$ and $\rho(f) \propto f^{-2}$ for $f \lesssim \gamma^{(n)}$ approximately. The scalings are obtained by Fourier transforming equations (1)–(4), applying the Wiener–Khintchine theorem, and converting frequency residuals to phase residuals (Meyers et al.

2021a, b; Antonelli et al. 2023; O’Neill, Meyers & Melatos 2024).⁶ The specific form of equations (1)–(4) is one particular choice of stochastic model for pulsar timing noise, whose PSD $\rho(f)$ is broadly consistent with observations for several MSPs, especially over typical inter-TOA intervals $T_{\text{cad}} = t_{k+1} - t_k \sim 1$ week. Alternative models exist, which include higher derivatives of $f_p^{(n)}(t)$ (Vargas & Melatos 2023) or multiple internal stellar components (Meyers et al. 2021a, b; Antonelli et al. 2023; O’Neill et al. 2024). The state-space framework presented in this paper extends straightforwardly to other stochastic models of pulsar spin wandering. Analogously, in traditional PTA analyses, whilst equation (B12) is a default description, multiple alternative parametrized models also exist (Sesana et al. 2008; Lentati et al. 2013; Sampson, Cornish & McWilliams 2015; Taylor, Simon & Sampson 2017; Chen et al. 2017a; Chen, Sesana & Del Pozzo 2017b; Aggarwal et al. 2019; Chen, Sesana & Conselice 2019). We refer the interested reader to Taylor (2021) for additional details on alternative power spectral density models.

APPENDIX C: SYNTHETIC PTA SPECIFICATIONS

The synthetic PTA deployed for testing in this paper is constructed to mimic the $N = 47$ pulsars from the 12.5-yr NANOGrav data release, whose sky positions are plotted in Fig. 2. The construction recipe is described in full in Section 4.1. In this appendix, we record a complete list of the pulsar parameters in Table C1, to assist the interested reader in implementing the analysis scheme and reproducing the test results in Sections 5–7.

APPENDIX D: DISPERSION OF θ_{gw} ESTIMATES

Every random realization of the noise processes $\xi^{(n)}(t)$ and $\varepsilon^{(n)}(t)$ leads to a different θ_{gw} posterior, when the synthetic data are analysed according to the procedure in Section 3.4. The distance between two posteriors (i.e. how similar they are) can be measured by many valid metrics, including those related to the Kolmogorov–Smirnov test (Corder & Foreman 2014). In this paper, we use the WD (Vaserstein 1969; Villani 2009) which is popular in machine learning (e.g. Arjovsky et al. 2017) and other domains. In Appendix D1, we define the WD and summarize its main properties. In Appendix D2, we present for reproducibility the WD calculated between every pair of posteriors for each static parameter in θ_{gw} across the 10^3 noise realizations of Section 5.4. The results are summarized via the median values reported in Table 2.

D1 Overview of the WD

The WD is a metric that defines a distance between two probability distributions $\mu(x)$ and $\nu(x)$. It has an intuitive interpretation as the lowest total cost with which one can move probability mass from μ to ν , with respect to a cost function $c(x, y)$. For this reason, it is sometimes known as the ‘Earth mover’s distance’. The p th order WD between two distributions is

$$W_p(\mu, \nu) = \left[\inf_{\gamma \in \Gamma(\mu, \nu)} \int c(x, y)^p d\lambda(x, y) \right]^{1/p}, \quad (\text{D1})$$

⁶The Wiener–Khintchine theorem assumes that the timing noise statistics are stationary, which may not be true in all pulsars. Non-stationarity is a subtle topic, which is deferred to future work.

Table C1. List of the fiducial pulsar parameters for the 47 pulsars used to construct the synthetic PTA in Section 4.1. The right ascension and declination of an individual pulsar in J2000 coordinates are labelled by α and δ , respectively. Parameter values are obtained from the ATNF pulsar catalogue (Manchester et al. 2005) using the PSRQPY package (Pitkin 2018).

JNAME	$f_{\text{em}}(t_1)$ (Hz)	$f_{\text{em}}^{(n)}(t_1) \times 10^{16}$ (s ⁻²)	α (deg)	δ (deg)	d (kpc)
J0023 + 0923	327.8	-12.3	5.8	9.4	1.8
J0030 + 0451	205.5	-4.3	7.6	4.9	0.3
J0340 + 4130	303.1	-6.5	55.1	41.5	1.6
J0613 - 0200	326.6	-10.2	93.4	-2.0	0.9
J0636 + 5128	348.6	-4.2	99.0	51.5	0.7
J0645 + 5158	112.9	-0.6	101.5	52.0	1.2
J0740 + 6620	346.5	-14.6	115.2	66.3	1.1
J0931 - 1902	215.6	-1.7	142.8	-19.0	3.7
J1012 + 5307	190.3	-6.2	153.1	53.1	0.7
J1024 - 0719	193.7	-7.0	156.2	-7.3	1.2
J1125 + 7819	238.0	-3.9	171.5	78.3	0.9
J1453 + 1902	172.6	-3.5	223.4	19.0	1.3
J1455 - 3330	125.2	-3.8	223.9	-33.5	0.7
J1600 - 3053	277.9	-7.3	240.2	-30.9	1.9
J1614 - 2230	317.4	-9.7	243.7	-22.5	0.7
J1640 + 2224	316.1	-2.8	250.1	22.4	1.5
J1643 - 1224	216.4	-8.6	250.9	-12.4	0.7
J1713 + 0747	218.8	-4.1	258.5	7.8	1.3
J1738 + 0333	170.9	-7.0	264.7	3.6	1.5
J1741 + 1351	266.9	-21.5	265.4	13.9	1.7
J1744 - 1134	245.4	-5.4	266.1	-11.6	0.4
J1747 - 4036	607.7	-48.5	267.0	-40.6	7.1
J1832 - 0836	367.8	-11.2	278.1	-8.6	2.1
J1853 + 1303	244.4	-5.2	283.5	13.1	2.1
J1857 + 0943	186.5	-6.2	284.4	9.7	1.2
J1903 + 0327	465.1	-40.7	285.8	3.5	7.0
J1909 - 3744	339.3	-16.1	287.4	-37.7	1.1
J1910 + 1256	200.7	-3.9	287.5	12.9	1.5
J1911 + 1347	216.2	-7.9	288.0	13.8	1.4
J1918 - 0642	130.8	-4.4	289.7	-6.7	1.1
J1923 + 2515	264.0	-6.7	290.8	25.3	1.2
J1939 + 2134	641.9	-433.1	294.9	21.6	3.5
J1944 + 0907	192.9	-6.4	296.0	9.1	1.2
J1946 + 3417	315.4	-3.1	296.6	34.3	6.9
J1955 + 2908	163.0	-7.9	298.9	29.1	6.3
J2010 - 1323	191.5	-1.8	302.7	-13.4	2.4
J2017 + 0603	345.3	-9.5	304.3	6.1	1.4
J2033 + 1734	168.1	-3.1	308.4	17.6	1.7
J2043 + 1711	420.2	-9.3	310.8	17.2	1.4
J2145 - 0750	62.3	-1.2	326.5	-7.8	0.7
J2214 + 3000	320.6	-15.1	333.7	30.0	0.6
J2229 + 2643	335.8	-1.7	337.5	26.7	1.8
J2234 + 0611	279.6	-9.4	338.6	6.2	1.0
J2234 + 0944	275.7	-15.3	338.7	9.7	1.6
J2302 + 4442	192.6	-5.1	345.7	44.7	0.9
J2317 + 1439	290.3	-2.0	349.3	14.7	1.7
J2322 + 2057	208.0	-4.2	350.6	21.0	1.0

for $p \in [1, \infty)$, where $\lambda(x, y)$ is the transport plan, and $\Gamma(\mu, \nu)$ is the set of all joint probability distributions for (x, y) that have marginals μ and ν , that is, $\Gamma(\mu, \nu)$ is the set of couplings of μ and ν .

The cost function can be freely chosen to suit the nature of the problem. Often, as in this paper, it is taken to be the absolute value function, $c(x, y) = |x - y|$. In general, $W_p(\mu, \nu)$ with $c(x, y) = |x - y|$ can be computed from n samples by the Hungarian algorithm (Kuhn 1955) in polynomial time $\mathcal{O}(n^3)$ (Villani 2009). However, for μ and ν defined on \mathbb{R}^d it is well-known (Dudley 1969) that $E[W_p(\mu, \nu)]$ converges slowly $\propto n^{-1/d}$. In this paper, we calculate

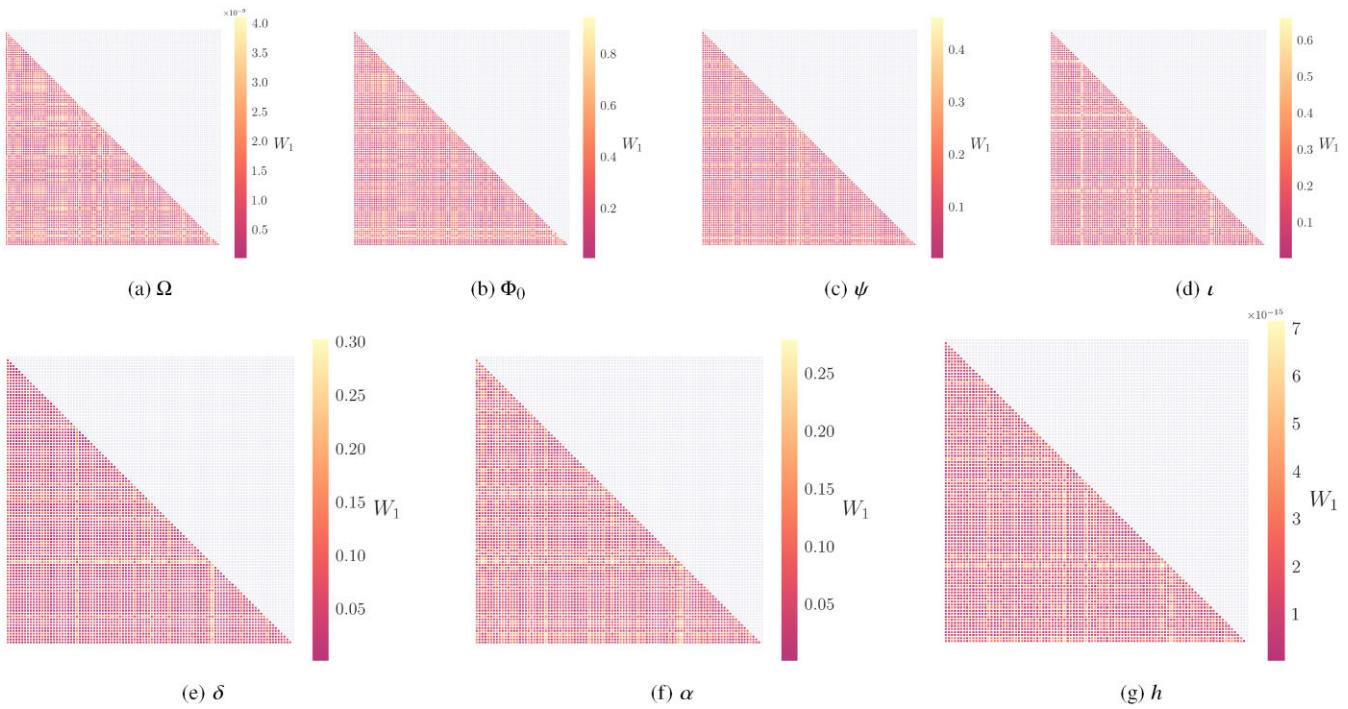


Figure D1. First moment of the WD, W_1 (colour scale), calculated between each pair of one-dimensional posteriors for the representative system in Table 1 across 10^3 realizations as discussed in Appendix D2, for each component of θ_{gw} . W_1 provides an upper bound on the difference in expected values between any two probability distributions. $W_1(\theta)$ is the vector for variable θ of $\binom{100}{2} W_1$ values that are plotted in each panel. The Pearson correlation coefficient between $W_1(\iota)$ and $W_1(h)$ (i.e. the ι - h panels) is 0.75, and between the $W_1(\delta)$ and $W_1(\alpha)$ (i.e. the δ - α panels) is 0.60.

the WD exclusively between the one-dimensional marginalized posteriors, setting $d = 1$ and obtaining

$$W_p(\mu, \nu) = \left\{ \int_0^1 dz [F_\mu^{-1}(z) - F_\nu^{-1}(z)] \right\}^{1/p}, \quad (\text{D2})$$

where $F_\mu^{-1}(z)$ is the inverse cumulative distribution function of μ .

The WD holds certain advantages over valid alternatives like the Kullback–Leibler divergence or the Kolmogorov–Smirnov test (Gelman et al. 2013; Corder & Foreman 2014). It is intuitive, being the minimum cost required to transform one distribution into another. It satisfies the Monge–Kantorovich duality (Villani 2003, 2009),

$$|E(X_\mu) - E(Y_\nu)| \leq W_1(\mu, \nu), \quad (\text{D3})$$

where X_μ and Y_ν are random variates drawn from the distributions μ and ν , respectively, and E denotes the expected value. It also satisfies desirable properties of a measure of distance, such as symmetry and the triangle inequality. For $p = 1$, the WD inherits the units of the underlying distributions. Finally, the WD is versatile; any two distributions can be compared, irrespective of whether they are continuous, discrete, or singular.

D2 WD for θ_{gw}

Table 2 summarizes the median value of the first moment of the WD, W_1 , between the 5×10^5 pairs of one-dimensional posteriors across the 10^3 realizations analysed in Section 5.4, for each of the seven parameters in θ_{gw} . To assist with reproducibility, we present in Fig. D1, the WD values for a subset (numbering 10^2) of the realizations. We plot a subset of realizations to avoid overcrowding the figure. The

subset corresponds to the 100 realizations plotted in Fig. 5. Fig. D1 contains seven subplots, one for each element of θ_{gw} . Each subplot is a lower triangular heat map, where each point denotes the W_1 value between a pair of one-dimensional posteriors for that element. Lower values of W_1 are magenta; higher values of W_1 are yellow. Note that the heat-map colour scale is not the same in every subplot. Instead the colour scale is set by the minimum and maximum W_1 values for each parameter. Normally it would be preferable to set the colour scale by the domain of the prior, but doing so renders the heat-map uniform, because the WD is generally much smaller than the prior domain (cf. Table 2).

Fig. D1 is consistent with the summary results presented in Table 2, and agrees with the conclusions drawn in Section 5.4, namely that the nested sampling scheme repeatedly converges to similar posteriors, for different realizations of the data, although dispersion remains. The W_1 value is small compared to the width of the prior, cf. Fig. 5 and $W_{1,\text{median,prior}}$ in Table 2. The W_1 value is significant compared to the injection value (up to 140 per cent for h_0 , typically $\lesssim 10$ per cent for other parameters; see Table 2), cf. $W_{1,\text{median,inj}}$ in Table 2. Qualitatively, Fig. D1 exhibits some correlated structure between panels. Specifically, Figs D1(d) and (g) (ι and h , respectively) appear correlated, as do Figs D1(e) and (f) (δ and α , respectively). Quantitatively, we define $W_1(\theta)$ as the vector of length $\binom{100}{2} W_1$ values for variable θ that are plotted in each panel of Fig. D1. The Pearson correlation coefficient between $W_1(\iota)$ and $W_1(h)$ is 0.75, and between the $W_1(\delta)$ and $W_1(\alpha)$ is 0.60.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.