

First results from the UTMOST-NS pulsar timing programme

L. Dunn^{1,2★}, C. Flynn^{2,3}, M. Bailes^{2,3}, Y. S. C. Lee^{1,2}, G. Howitt^{1,2,4,5}, A. Melatos^{1,2}, V. Gupta⁶,
A. Mandlik³ and A. Deller³

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

²*Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Hawthorn, VIC 3122, Australia*

³*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Mail H30, PO Box 218, Hawthorn, VIC 3122, Australia*

⁴*Computational Biology Program, Peter MacCallum Cancer Centre, Parkville, VIC 3010, Australia*

⁵*Sir Peter MacCallum Department of Oncology, University of Melbourne, Parkville, VIC 3010, Australia*

⁶*Australia Telescope National Facility, CSIRO, Space and Astronomy, PO Box 76, Epping, NSW 1710, Australia*

Accepted 2025 June 27. Received 2025 June 5; in original form 2024 August 26

ABSTRACT

The UTMOST-NS pulsar timing programme operated at the Molonglo Observatory Synthesis Telescope from 2021 April to 2023 June, observing 173 pulsars with an average cadence of 50 pulsars per day. An overview of the programme is presented, detailing the hardware, software, and observing strategy. Pulsar timing results are discussed, focusing on timing noise and glitches. It is shown that the scaling of residuals due to timing noise with pulsar parameters and observing time span is consistent with earlier studies, and that the recovered timing noise parameters remain consistent as the observing time span is increased. Second frequency derivatives are investigated, and it is shown that the uncertainty on $\ddot{\nu}$ is sensitive to the frequency cut-off in the timing noise model, varying by threefold approximately depending on whether Fourier modes with frequency lower than the reciprocal of the observing time span are included. We measure 39 non-zero values of $\ddot{\nu}$ when considering both models with and without low-frequency modes. An analytic scaling relating anomalous braking indices to timing noise amplitude is also validated. Glitches in the sample are discussed, including three detected by an ‘online’ glitch detection pipeline using a hidden Markov model (HMM). In total 17 glitches are discussed, one of which, in PSR J1902+0615, has not been reported elsewhere. An ‘offline’ glitch search pipeline using the HMM framework is used to search for previously undetected glitches. Systematic upper limits are set on the size of undetected glitches. The mean upper limit is $\Delta v^{90\%}/\nu = 6.3 \times 10^{-9}$ at 90 per cent confidence.

Key words: stars: neutron – pulsars: general – stars: rotation.

1 INTRODUCTION

Precision studies of the rotational behaviour of slow (spin period $\gtrsim 0.1$ s) radio pulsars can provide a wealth of information on the extreme environments in and around neutron stars. In particular, rotational irregularities known as ‘timing noise’ and ‘glitches’ can be used to constrain the composition and internal structure of these objects, and the nature of the coupling between the various components (Haskell & Melatos 2015; Chamel 2017).

Timing noise appears as a persistent, time-correlated stochastic wandering in the phase of the pulsar, with amplitude varying over many orders of magnitude throughout the pulsar population. The details differ from pulsar to pulsar, in some cases being modelled as a random walk in the phase, frequency or frequency derivative (Boynton et al. 1972; Cordes & Helfand 1980), while in other cases, it appears that a random walk is unable to account for the observed behaviour (Cordes & Downs 1985; D’Alessandro et al. 1995; Hobbs, Lyne & Kramer 2010). Timing noise is frequently decomposed as a sum of sinusoids with a power-law spectrum, possibly with a corner frequency or a high-frequency cut-off (Lentati et al. 2013; Parthasarathy et al. 2020; Reardon et al. 2023). Phenomenological

modelling of this type does not directly address what causes the behaviour in the first place, although a number of mechanisms have been proposed, often involving a superfluid component in the interior of the star (Alpar, Nandkumar & Pines 1986; Jones 1990; Melatos & Link 2014; Meyers, Melatos & O’Neill 2021).

Timing noise is an important confounding factor for performing timing measurements of slow pulsars. It primarily arises in the context of measuring the second frequency derivative $\ddot{\nu}$, which is of particular physical interest because of its connection with the long-term braking mechanisms at play. Many physical braking mechanisms imply a relation of the form $\dot{\nu} = K\nu^n$, where n is referred to as the ‘braking index’ and can be measured (assuming K is constant) by the combination $n = \nu\ddot{\nu}/\dot{\nu}^2$ (Lyne & Graham-Smith 2012). Values of $\ddot{\nu}$ derived from analyses which are not timing noise-aware are often contaminated (Hobbs et al. 2010; Chukwude & Chidi Odo 2016; Onuchukwu & Legahara 2024) – such values of $\ddot{\nu}$ are in some cases themselves used as a measure of timing noise strength (Taylor 1991; Matsakis, Taylor & Eubanks 1997). However, under certain circumstances even analyses which do incorporate a timing noise model can fail to fully account for the influence of timing noise on the measured value of $\ddot{\nu}$ (Keith & Nişu 2023; Vargas & Melatos 2023).

Glitches, in contrast to the continuous perturbations of timing noise, are step changes in the pulse frequency. They are often

* E-mail: liam@ldunn.me

associated with a step change in frequency derivative, and an exponential relaxation back towards the pre-glitch trend (Lyne & Graham-Smith 2012). This phenomenological model of a glitch and its recovery is widely used, but as with timing noise the question of the underlying physical mechanism is far from settled (see e.g. Haskell & Melatos 2015, for a review). Often the superfluid interior is invoked, as significant angular momentum may be transferred from this component to the solid crust, to which the electromagnetic emission is tied (Alpar et al. 1984; Link & Epstein 1991; Jones 1998; Peralta et al. 2006). It has also been suggested that mechanical failure of the crust under strain (a ‘starquake’) may play an important role in the glitch mechanism (Ruderman 1976; Link & Epstein 1996; Middleditch et al. 2006; Kerin & Melatos 2022).

The mechanisms involved in the glitch phenomenon may be probed both by close study of individual glitches, when sufficiently detailed data are available (Dodson, McCulloch & Lewis 2002; Palfreyman et al. 2018; Ashton et al. 2019a), and population-level studies investigating both the statistics of glitches in individual pulsars (Melatos, Peralta & Wyithe 2008; Espinoza et al. 2014; Melatos, Howitt & Fulgenzi 2018; Carlin & Melatos 2019; Carlin, Melatos & Vukcevic 2019; Fuentes, Espinoza & Reisenegger 2019; Carlin & Melatos 2020; Ho et al. 2020) as well as the statistics across the pulsar population as a whole (Lyne, Shemar & Smith 2000; Espinoza et al. 2011; Fuentes et al. 2017; Eya, Urama & Chukwude 2019; Lower et al. 2021; Millhouse et al. 2022). The latter especially are reliant on large-scale observing campaigns (Espinoza et al. 2011; Yu et al. 2013; Jankowski et al. 2019; Basu et al. 2022; Keith et al. 2024) and reliable glitch detection strategies (Espinoza et al. 2014; Yu & Liu 2017; Melatos et al. 2020; Singha et al. 2021a).

In pursuit of a data set which could form a platform for studies of these phenomena, the UTMOST-NS pulsar timing programme operated from 2021 April to 2023 June, monitoring 173 pulsars with daily to weekly cadence. The UTMOST-NS programme used the refurbished north–south arm of the Molonglo Observatory Synthesis Telescope (MOST), and was undertaken as part of the broader UTMOST project (Bailes et al. 2017; Deller & Flynn 2020; Day 2022; Mandlik et al. 2024) which had already revived the east–west arm as a pulsar timing instrument (hereafter referred to as UTMOST-EW) (Jankowski et al. 2019; Lower et al. 2020). In large part, the UTMOST-NS programme was a continuation of the UTMOST-EW timing effort, with similar motivation and significant overlap in targets.

The paper begins with a discussion of the observational foundations of the timing programme (Section 2), from the details of the telescope configuration to the observing strategy and the radio frequency interference (RFI) environment. We then set out the pulsar timing framework which is used throughout the rest of this work (Section 3), before moving on to investigations of timing noise (Section 4), second frequency derivatives (Section 5), and glitches (Section 6) in the context of the UTMOST-NS data set. We conclude and look to future work in Section 7.

2 OBSERVATIONS

2.1 System overview

The pulsar timing programme is performed using the north–south arm of the MOST, which was refurbished as part of the UTMOST¹ project (Bailes et al. 2017; Deller & Flynn 2020; Day 2022). The refurbished system is referred to as UTMOST-NS. A complete

description of the system can be found in Mandlik et al. (2024); here we give a brief overview of its basic characteristics.

The north–south arm of the MOST is composed of two 778 m × 12.7 m paraboloid cylindrical reflectors, with a total collecting area of 19800 m². UTMOST-NS consists of 66 × 1.4 m long individual array elements called cassettes, each containing eight dual-pole patch antennas, which are primarily installed in a ‘dense core’ near the centre of the arm. The total collecting area visible to these cassettes is 1375 m².

The operating frequency range of UTMOST-NS is 810–860 MHz. The typical system equivalent flux density (SEFD) for the timing measurements is 680 Jy, and the typical system temperature is 170 K. The performance of the telescope over the duration of the programme and the estimation of flux densities for pulsars in our sample are discussed in Appendix A.

UTMOST-NS operates as a transit instrument with no moving parts – pointing along the meridian is performed electronically via beamformers. The primary beam has an extent of 2.5° in the east–west direction, 12.7° in the north–south direction, and the beam can be pointed in the north–south direction from –90° to approximately +25° in declination. Pulsars are thus typically visible within the half-power points of the primary beam for between 5 and 25 min, depending on their declination.

Sky frequencies are downconverted using a 900 MHz signal, which is referenced to a station clock, consisting of a Valon 5007 frequency synthesizer disciplined by a 10 MHz GPS standard.

We primarily use the bright millisecond pulsar (MSP) PSR J0437–4715 to monitor and correct for small clock jumps, which typically take place after losses of power to the entire system, or other disruptions. Over the course of the timing program, six such corrections are logged, lying in the range 4.2–18 μs. These corrections are cross-checked for consistency against another MSP, J2241–5236. This low dispersion measure (DM) pulsar scintillates and is used as a consistency check on the clock rather than as a primary reference because it was not always detected in daily transits.

For timing purposes, we take the geographical reference position of the telescope to be the same position used for the pulsar timing programme on the east–west arm: a latitude of –35° 22′ 14.5518″, a longitude of 149° 25′ 28.8906″, and an elevation of 741 m in the International Terrestrial Reference Frame (2008) coordinates assuming a Geodetic Reference System 1980 ellipsoid. The latitude and longitude were determined in a geodetic survey made by Geoscience Australia in 2012 November, and the altitude was measured using a consumer GPS device (Garthwaite et al. 2013; Jankowski et al. 2019).

2.2 Target selection and observing strategy

The primary science goal for the UTMOST-NS timing programme is the study of timing irregularities in slow pulsars. The target selection and observing strategy is focused on high-cadence (typically ≲ 5 d) monitoring of those slow pulsars which have high enough mean flux densities (≳ 5 mJy at 840 MHz, see Appendix A) to be seen in the time they take to transit the primary beam.

An initial list of approximately 250 candidate pulsars was selected using the mean flux densities in the timing program on the east–west arm of the array (Jankowski et al. 2019). Timing data were acquired for this set of pulsars during commissioning of the upgrade to the north–south arm in 2020 and 2021, and the timing list pared down to 173 pulsars to be monitored regularly based on whether the pulsars were successfully detected in these early observations.

Although the majority of the target pulsars are slow, we include a handful of MSPs, primarily for the purpose of monitoring the

¹UTMOST is not an acronym.




Figure 1. Sky locations of the 173 pulsars observed in the UTMOST-NS survey, shown in RA in hours and Dec. in degrees. The symbol size is proportional to mean flux density. Vertical grey zones show groupings of pulsars based on proximity in RA. A single black symbol in such a zone indicates a pulsar which can be timed daily; green/blue symbols show two pulsars which are timed on alternative days; green/blue/red shows pulsars timed every third day on a three day cycle, up to a maximum of six pulsars timed over six transit days in busy regions of the Galactic plane (particularly between 15 and 20 h RA, where the plane lies approximately north–south on the sky and there are many pulsars transiting the meridian). The horizontal width of coloured boxes containing each pulsar indicates the integration time (and is on the same scale as the RA axis, i.e. 1 h in RA is 1 h of elapsed time).

stability of the system clock. These are also particularly useful, because high-cadence monitoring of MSPs serves as a cross-check of timing integrity, when glitches or other step-changes in behaviour are found by our real-time detection system, such as the transient profile event observed in PSR J1737+0747 (Xu et al. 2021; Singha et al. 2021b; Jennings, Nanograv Collaboration & Chime/Pulsar Collaboration 2022). We also include the radio-loud magnetars PSR J1622–4950 (Levin et al. 2010) and PSR J1809–1943 (Camilo et al. 2006; Levin et al. 2019). Although we are not able to establish timing solutions for either magnetar due to low numbers of detections, PSR J1809–1943 is particularly valuable as a testbed for the fast radio burst (FRB) search pipeline ‘Adjacent Beam Classifier’ (ABC) developed for commensal operation with pulsar timing (Mandlik 2024).

The observing schedule is constructed semi-automatically. Many targets can be observed daily. In cases where multiple pulsars overlap in right ascension and the available transit time does not permit us to dwell on each for long enough to create a high-quality time of arrival (ToA), we cycle through sets of targets on consecutive days. Broadly speaking, pulsars in regions away from the Galactic plane are mostly observed daily, while pulsars in busy regions of the Galactic plane are observed with a 2- to 6-d cadence. Fig. 1 shows the locations of the 173 pulsars, and indicates regions of the sky in which daily, and 2- to 6-d cadences are performed by cycling through the available pulsars in transiting regions of sky. The observations are conducted automatically around the clock, based on these schedules, using the HADES observing system. ABC searches for FRBs both during pulsar

observations, and while the scheduler is waiting for the next program pulsar to transit the beam (Mandlik et al. 2024).

2.3 RFI environment

As with all radio facilities, contamination of pulsar observations by RFI presented challenges over the course of the timing programme. In particular, the UTMOST-NS operational frequency band (810–860 MHz) overlaps with handset emissions from two Australian mobile phone networks, as well as local sources of narrow-band interference (aka ‘birdies’). Very strong signals from 3G and 4G mobile phone base stations occur at 780–790 and 870–890 MHz, on either side of our operating band, requiring sharp roll-off high- and low-pass filters at the band edges. Due to the site location, which is a flat valley surrounded by hills, and the low local population density, these transmissions rarely rendered an entire observation unusable. Rather, they are well localized in time (typically lasting less than a minute) and frequency (occupying at most 20 per cent of the total observing band). This typical pattern of RFI is cleaned very successfully using XPROF, a tool which uses knowledge of the pulsar profile, flux density, and ephemeris to judiciously clean off-pulse regions in subintegration and frequency channel space.

Fig. 2 shows an example observation contaminated by cell network signals before and after cleaning. The frequency bands occupied by two cell network transmissions are clearly visible in the pre-cleaning frequency-phase plot (top left), occupying roughly 825–830 and 840–845 MHz. Each transmission occupies only one or




Figure 2. Before (top row) and after (bottom row) example demonstrating successful RFI cleaning using XPROF for the observation of PSR J1534–5334 taken at 2021-10-08-04:29:08, with the frequency-phase (left) and time-phase (middle) heat maps, and the fully integrated pulse profile (right). Significant contamination from multiple cell network transmissions (825–830 and 840–845 MHz) is visible in the pre-cleaning plots, and largely absent in the post-cleaning plots.

two 20 s subintegrations, as shown in the pre-cleaning time-phase plot (top middle). The rightmost plots show the fully integrated profiles pre- and post-cleaning. Although the pulse profile is clear in both plots, the pre-cleaning profile shows obvious contamination in the off-pulse region, which is largely removed by the cleaning procedure.

For the pulsars that were regularly timed in the programme, the overall median percentage of data deleted in an observation is 16 per cent. This includes the automatic 10 per cent deleted at the edges of the frequency band, as these regions are often overwhelmed by loud RFI, which interferes with XPROF’s ability to properly clean the data. The per-pulsar medians range between 11 per cent and 24 per cent.

2.4 Timing pipeline

There are several intermediate stages between the data stream acquired at the telescope and the pulse times of arrival which are the primary final data product of the timing programme.

The initial folded archives are produced from tied-array beam data at $10.24 \mu\text{s}$ time resolution using DSPSR (van Straten & Bailes 2011) which performed the coherent dedispersion and folding according to a stored ephemeris. These archives have 512 frequency channels covering the 50 MHz bandwidth and two orthogonal linear polarizations, although only the total intensity has been used in this work.

Once the raw-folded archives have been produced, they are passed through the RFI mitigation step – see Section 2.3 for a description of the RFI environment at Molonglo and the mitigation procedures.

The RFI-cleaned archives are finally processed into ToAs using PAT, a component of PSRCHIVE (Hotan, van Straten & Manchester 2004), to cross-correlate the observed pulse profile against the standard profile and thus derive a pulse arrival time and associated uncertainty. The standard profiles are generated by smoothing high signal-to-noise ratio (S/N) observations of each pulsar. Where available, we use the standard profiles generated for the UTMOST-EW pulsar timing programme (Jankowski et al. 2019; Lower et al. 2020). All 512 frequency channels of the archives are averaged together before




Figure 3. Timing residuals from UTMOST-NS observations of the MSP PSR J2241–5236. The weighted rms is $3.2 \mu\text{s}$.

cross-correlation with the profile templates. For pulsars which were not observed on the east–west arm, we obtained standard profiles from observations taken by the MeerKAT telescope (Johnston et al. 2020), or from UTMOST-NS observations. The generated ToAs and uncertainties are stored along with additional metadata (e.g. S/N for the observation, and duration of the observation). Only ToAs with S/N greater than 7 are regarded as ‘good’ ToAs to be used in further timing analysis. In total 18 360 such ToAs pass this cut out of the final data set of 34 575 UTMOST-NS observations. 370 observations which produce a ToA with S/N > 7 were judged to be RFI-contaminated and manually excluded, leaving 17 990 science-ready ToAs.

As a demonstration of the timing accuracy achieved by the UTMOST-NS system, Fig. 3 shows the timing residuals for the UTMOST-NS observations of the MSP PSR J2241–5236, covering the time period between 2021 May and 2023 June and comprising 120 times of arrival. The mean ToA uncertainty is $4.8 \mu\text{s}$, and the weighted rms of the residuals over the full observation period is $3.2 \mu\text{s}$.

3 PULSAR TIMING

In this section, we introduce the primary pulsar timing methodology used throughout this work. Section 3.1 presents the components of the timing model, including both deterministic and stochastic contributions, and Section 3.2 discusses the Bayesian model selection procedure used to decide which components should be incorporated into the timing model for a data set.

3.1 Timing framework

We adopt a pulsar timing framework which uses the ENTERPRISE (Ellis et al. 2020) pulsar inference package in conjunction with LIBSTEMPO (Vallisneri 2020) and TEMPO2 (Hobbs, Edwards & Manchester 2006). ENTERPRISE allows the user to construct a model for the pulsar ToAs which incorporates both deterministic and stochastic processes, i.e. $\mathbf{t} = \mathbf{t}_{\text{det}} + \mathbf{t}_{\text{sto}}$, where \mathbf{t} is the vector of ToAs. The deterministic contribution is generated by the standard phase model provided by TEMPO2 (Edwards, Hobbs & Manchester 2006), and contains both the intrinsic deterministic evolution of the pulsar, expressed as a low-order Taylor series in the time coordinate of the pulsar’s reference frame, as well as the kinematic Doppler contribution associated with the transformation between the pulsar

reference frame and the reference frame of the telescope. While there are in principle a large number of parameters which determine the form of \mathbf{t}_{det} , in this work we primarily restrict our attention to the pulsar’s intrinsic rotational frequency and its first two derivatives, ν , $\dot{\nu}$, and $\ddot{\nu}$, as well as the pulsar’s position on the sky, expressed in terms of right ascension α and declination δ .² In some cases further parameters are necessary to obtain a good fit, e.g. binary orbital elements and proper motions, but we leave these fixed at the values recorded in the first UTMOST data release (Lower et al. 2020) where applicable, or the Australia Telescope National Facility pulsar catalogue (Manchester et al. 2005). In cases where a clear glitch has been detected from visual inspection of the residuals we include a corresponding glitch term in our phase model, viz.

$$\phi_g(t) = \left[\Delta\phi_p + \Delta\nu_p(t - t_g) + \frac{1}{2}\Delta\dot{\nu}_p(t - t_g)^2 - \tau_d\Delta\nu_d e^{-(t-t_g)/\tau_d} \right] \Theta(t - t_g). \quad (1)$$

Here, $\Delta\phi_p$, $\Delta\nu_p$, and $\Delta\dot{\nu}_p$ denote permanent step changes in ϕ , ν , and $\dot{\nu}$, respectively. The exponential term represents a decaying step in frequency of magnitude $\Delta\nu_d$ which relaxes on a time-scale τ_d . The factor $\Theta(t - t_g)$ is the Heaviside step function, and t_g indicates the glitch epoch. All parameters except t_g are allowed to vary when inferring the glitch properties.

There are two stochastic processes of interest. The first is white noise in the residuals due to ToA measurement uncertainty σ_{ToA} , calculated by PAT, and two additional parameters EFAC and EQUAD, which account for misestimation of σ_{ToA} due to e.g. non-ideal profile templates, unaccounted for instrumental effects, or intrinsic profile variation (Liu et al. 2011; Shannon et al. 2014), viz.

$$\sigma_{\text{out}}^2 = (\text{EFAC})^2 \sigma_{\text{ToA}}^2 + (\text{EQUAD})^2. \quad (2)$$

The second stochastic process is a long-term ‘wandering’ in the residuals, commonly known as timing noise. The contribution to \mathbf{t} from the latter process is assumed to have a power spectral density (PSD) of the form

$$P(f) = \frac{A_{\text{red}}^2}{12\pi^2 f_{\text{yr}}^3} \left(\frac{f}{f_{\text{yr}}} \right)^{-\beta}, \quad (3)$$

where A_{red} is the dimensionless amplitude of the timing noise process, β is the spectral index of the timing noise PSD, and $f_{\text{yr}} = 1/(1 \text{ yr})$ is an arbitrarily chosen reference frequency. This process is modelled in ENTERPRISE as a set of harmonically related sinusoids. By default the sinusoids have their lowest frequency at $1/T_{\text{span}}$, where T_{span} is the total length of the data set in question, and extend up to $30/T_{\text{span}}$ in increments of $1/T_{\text{span}}$. However, this is easily modified, and we also consider a ‘long-period’ timing noise model where the lowest frequency is instead at $1/(2T_{\text{span}})$, but still extending up to $30/T_{\text{span}}$ in increments of $1/(2T_{\text{span}})$, so this model has twice as many sinusoids overall. Although it is possible to fit for the amplitude of each harmonic separately and then combine these into an estimate of the form of the PSD, here we opt instead to constrain the PSD to take the form of equation (3) and estimate A_{red} and β . In this approach, A_{red} and β play the role of hyperparameters: the prior distributions on the amplitudes of each harmonic are fixed by $P(f)$, and subsequently marginalized over (Lentati et al. 2013). We do not include models for chromatic stochastic processes, such as variations in the DM or scattering (Goncharov et al. 2021). Given

²We opt to sample all of the timing model parameters rather than marginalize over them – this is motivated primarily by ease of implementation.

the small fractional bandwidth of our observations, these effects are covariant with achromatic red noise (Lower et al. 2020).

We adopt a Bayesian approach to inferring the various parameters of the different processes in the pulsar timing models. ENTERPRISE constructs an appropriate likelihood function for a model \mathcal{M} , $\mathcal{L}(\mathbf{t} | \boldsymbol{\theta}, \mathcal{M})$, where $\boldsymbol{\theta}$ is a vector of parameter values. Given this likelihood along with a prior distribution on the parameters of interest, $\Pr(\boldsymbol{\theta})$ (see table 1), we use the MULTINEST (Feroz, Hobson & Bridges 2009) nested sampling package to explore the parameter space and estimate the posterior distribution of $\boldsymbol{\theta}$,

$$\Pr(\boldsymbol{\theta} | \mathbf{t}, \mathcal{M}) = \frac{\mathcal{L}(\mathbf{t} | \boldsymbol{\theta}, \mathcal{M}) \Pr(\boldsymbol{\theta})}{\Pr(\mathbf{t} | \mathcal{M})} \quad (4)$$

as well as the model evidence,

$$\Pr(\mathbf{t} | \mathcal{M}) = \int d\boldsymbol{\theta} \mathcal{L}(\mathbf{t} | \boldsymbol{\theta}, \mathcal{M}) \Pr(\boldsymbol{\theta}). \quad (5)$$

The model evidence is also denoted by \mathcal{Z} .

3.2 Parameter inference and model selection

To go from the pulsar timing framework described in the previous section to inference results for the pulsars in our sample, we must first decide on what data are to be incorporated into the analyses. For those pulsars which were timed as part of both the UTMOST-EW and UTMOST-NS timing programmes, we combine the two data sets to maximize the timing baseline. This includes a majority of the UTMOST-NS pulsars, with 93 per cent (150) having previous data from the UTMOST-EW programme, and 83 per cent (134) with a total timing baseline of more than 7 yr. Those pulsars for which only UTMOST-NS data exist have timing baselines between approximately 6 months and 2 yr.

The values of the timing model parameters are informed by an initial least-squares fit performed using TEMPO2, which for each timing model parameter p returns a best-fitting value p^* and an associated error estimate Δ_p . The priors on the timing model parameters are then taken to be uniform on $p^* \pm 5000\Delta_p$, with the following exceptions. When performing the initial least-squares fit we do not fit for $\dot{\nu}$, and the prior on $\dot{\nu}$ is taken to be $\pm 5000\sigma_{\text{rms}} T_{\text{span}}^{-3}$. If a glitch is included in the fit, then the initial least-squares fit is performed with no decay term included. The parameters $\Delta\phi$ and $\Delta\dot{\nu}_p$ are treated in the same way as other timing model parameters. However, the permanent and decaying frequency increment are taken to have a log-uniform distribution on $[10^{-10} \text{ Hz}, 10^{-4} \text{ Hz}]$, while the decay time-scale is taken to have a log-uniform distribution on $[10^{-1} \text{ d}, 10^3 \text{ d}]$. The log-uniform distributions on the frequency increments reflect the fact that a priori their individual orders of magnitude are not well constrained, even though their sum is reasonably well measured by the initial least-squares fit. See Section 6.5 for further discussion on the treatment of the decaying component in the glitch parameter estimation.

For each pulsar, we consider two possible models. One is TNF2, which comprises a deterministic timing model including a $\dot{\nu}$ term, and timing noise modelled by harmonically related sinusoids parametrized by A_{red} and β [see equation (3)] with a low-frequency cut-off $1/T_{\text{span}}$. The other is TNLONGF2, which is identical to TNF2 except that the low-frequency cut-off of the timing noise model is $1/(2T_{\text{span}})$ instead. We compute the model evidence \mathcal{Z} for both models using MULTINEST, and evaluate the Bayes factor $\mathcal{B} = \mathcal{Z}_{\text{TNLONGF2}}/\mathcal{Z}_{\text{TNF2}}$. If we obtain $\ln \mathcal{B} > 5$, we adopt TNLONGF2 for that pulsar in subsequent analyses, otherwise we adopt TNF2.

4 TIMING NOISE

A large fraction of the pulsars in our sample exhibit detectable levels of timing noise, and in this section, we investigate its properties across the sample.

In order to compare timing noise between data sets, it is useful to adopt some standard measure of timing noise strength which gives an indication of the contribution from the modelled timing noise to the timing residuals. Here, we follow Parthasarathy et al. (2019) and Lower et al. (2020) and adopt the following figure of merit, based on the values of the timing noise parameters returned by ENTERPRISE [see equation (3)] and the low-frequency cut-off in the timing noise model, f_{low} :

$$\sigma_{\text{RN}}^2 = \int_{f_{\text{low}}}^{\infty} P(f) df \quad (6)$$

$$= f_{\text{yr}}^{-3+\beta} \frac{A_{\text{red}}^2}{12\pi^2} \frac{f_{\text{low}}^{-(\beta-1)}}{\beta-1}. \quad (7)$$

We set $f_{\text{low}} = 1/T_{\text{span}}$ for TNF2, and $1/2T_{\text{span}}$ for TNLONGF2.

In total, there are 89 pulsars in our sample for which σ_{RN} is larger than $3\sigma_{\text{ToA}}$, where σ_{ToA} is the mean ToA uncertainty for that pulsar in our data set. These are the pulsars which we take as having ‘significant’ timing noise.

4.1 Correlation between timing noise activity and spin parameters

A number of authors have investigated how timing noise strength correlates with pulsar spin parameters (e.g. Cordes & Helfand 1980; Dewey & Cordes 1989; Urama, Link & Weisberg 2006; Hobbs et al. 2010; Shannon & Cordes 2010; Melatos & Link 2014; Parthasarathy et al. 2019; Lower et al. 2020), generally finding that there is a significant correlation of timing noise activity with $\dot{\nu}$, for a variety of activity measures and a variety of methodologies for assessing the correlations. Here, following Dewey & Cordes (1989), we look to estimate σ_{RN} based on the spin frequency and frequency derivative, i.e. we construct

$$\chi_{\text{RN}} = C \nu^a |\dot{\nu}|^b f_{\text{low}}^{-\gamma} \quad (8)$$

and aim to estimate C , a , b , and γ such that χ_{RN} accurately estimates σ_{RN} (χ_{RN} has units of seconds, as does σ_{RN}).

As in Shannon & Cordes (2010), we assume that the residuals between χ_{RN} and σ_{RN} are lognormally distributed and hence adopt the likelihood

$$\mathcal{L} = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\epsilon}} \exp \left\{ -\frac{[\log_{10}(\sigma_{\text{RN},i}) - \log_{10}(\chi_{\text{RN},i})]^2}{2\epsilon^2} \right\}, \quad (9)$$

where the product is over all N pulsars with significant timing noise, indexed by i , and ϵ is an additional parameter which describes the scatter in the relation between σ_{RN} and χ_{RN} . The scatter in the relation may arise from random variation in the timing noise realizations, or from variations in other parameters which are not included in χ_{RN} [e.g. neutron star composition and temperature (Alpar et al. 1986; Melatos & Link 2014)], or it may arise from a breakdown in the simplifying assumption that the values of C , a , b , and γ are universal across the population. For this investigation we do not treat these potentially complex issues, and simply absorb all of these sources of variation into ϵ , except that the question of variation in γ is discussed at the end of this section.

We take (log-)uniform priors, with a , b , γ , $\log_{10} C \in [-10, 10]$ and $\epsilon \in [0, 10]$. We use BILBY (Ashton et al. 2019b) in conjunction with the DYNESTY (Speagle 2020) nested sampling package to obtain

Table 1. Prior ranges used in ENTERPRISE analyses described in Section 3. All prior distributions are (log-)uniform. For a timing model parameter p , the best-fitting value of the parameter returned by an initial TEMPO2 fit is denoted by p^* , and the associated uncertainty is denoted by Δp . The rms of the phase residuals from the TEMPO2 fit is denoted by σ_{rms} .

Parameter	Prior range
$\alpha, \delta, \nu, \dot{\nu}$	$p^* \pm 5000\Delta p$
$\ddot{\nu}$	$\pm 5000\sigma_{\text{rms}}T_{\text{span}}^{-3}$
$\Delta\phi, \Delta\dot{\nu}_p$	$p^* \pm 5000\Delta p$
$\log_{10}(\Delta\nu_p/\text{Hz}), \log_{10}(\Delta\nu_d/\text{Hz})$	$[-10, -4]$
$\log_{10}(\tau_d/d)$	$[-1, 3]$
$\log_{10}(\text{EFAC})$	$[-1, 0.7]$
$\log_{10}(\text{EQUAD}/s)$	$[-8, -2]$
$\log_{10}(A_{\text{red}}/\text{yr}^{3/2})$	$[-14, -7]$
β	$[2, 10]$

posterior distributions on all parameters – the parameter estimates are summarized in the first column of Table 2.

The estimated values of the a , b , and γ parameters agree broadly with the earlier studies of Shannon & Cordes (2010) and Lower et al. (2020) [although we note that our data sets are not independent of Lower et al. (2020), as we include a subset of the UTMOST-EW data which were used in that analysis]. As in those cases, we find $a \approx -b$ within the quoted uncertainties, which are comparable; we find roughly ± 0.4 for a and ± 0.2 for b , noting that Shannon & Cordes (2010) quote 2σ confidence intervals, while we quote 95 per cent credible intervals, as do Lower et al. (2020). This suggests that the characteristic age $\tau_c \sim \nu|\dot{\nu}|^{-1}$ is a good proxy for σ_{RN} , at least for slow pulsars, which make up a majority of this sample. It has been noted that τ_c is also a good proxy for a pulsar’s glitch rate, which may or may not involve related physical processes. Millhouse et al. (2022) showed that a model in which glitch rate depends on τ_c is preferred over a model where the glitch depends on a generic combination $\nu^a \dot{\nu}^b$ with $b \neq a$, or ν^a or $\dot{\nu}^b$ alone. We also find agreement in the value of γ , although there is more variation in both the point estimate and the uncertainties – our value of $\gamma \approx 1.2 \pm 0.8$ lies between the Lower et al. (2020) value of 1.0 ± 1.2 and the Shannon & Cordes (2010) value of 1.9 ± 0.2 . The degree of scatter is also similar among all three analyses, with $\epsilon \approx 0.7$ in all cases.

For simplicity, we assume so far that γ takes one value across the population, but the form of σ_{RN} [equation (7)] suggests that if A_{red} is constant over the observing span we should expect $\gamma = (\beta - 1)/2$. We therefore obtain a second set of estimates for C , a , b , and ϵ , using the same functional form as equation (8) but with $\chi_{\text{RN}} \propto f_{\text{low}}^{-(\beta-1)/2}$, where β is taken to be the mean posterior value from the favoured

model obtained via analysis with ENTERPRISE. The results are given in the second column of Table 2. The a and b values are significantly different from the $f_{\text{low}}^{-\gamma}$ analysis – in particular the value of a is significantly lower, and more consistent with $a = -b$. We note that in this case our posteriors are roughly centred around $a = -b = -0.5$, whereas the Shannon & Cordes (2010) and Lower et al. (2020) values are more consistent with $a = -b = -1$. The scatter is reduced relative to the $\chi_{\text{RN}} \propto f_{\text{low}}^{-\gamma}$ model, as indicated by the ϵ values, and the $f_{\text{low}}^{-(\beta-1)/2}$ model is favoured over the $f_{\text{low}}^{-\gamma}$ model by a log Bayes factor of 11.

The model adopted here is of course highly simplified. There are many other pulsar properties which may influence the timing noise behaviour and are not included in our basic model, e.g. magnetic field structure, internal composition, and temperature (Pons, Viganò & Geppert 2012; Melatos & Link 2014; Haskell & Sedrakian 2018). Variation in these properties is likely to account at least partially for the observed scatter in the relation between σ_{RN} and χ_{RN} . In the particular case of the $\chi_{\text{RN}} \propto f_{\text{low}}^{-\gamma}$ results, it is plausible that the higher scatter relative to the Lower et al. (2020) results is due to the wider range of T_{span} in our sample – we have $1.9 < T_{\text{span}}/\text{yr} < 8.7$, whereas the Lower et al. (2020) sample has $1.4 < T_{\text{span}}/\text{yr} < 4.7$. The decrease in scatter when the spectral index of the timing noise is incorporated explicitly into the f_{low} scaling similarly points towards the inadequacy of a model with a universal f_{low} scaling relation.

4.2 Consistency of UTMOST-EW and -NS timing noise parameter estimates

In the previous section, we noted that our results are not independent of Lower et al. (2020), as our data set includes a subset of the full UTMOST-EW data set analysed in that work, augmented with the addition of the UTMOST-NS data which extends the time span covered. These two data sets are analysed with similar, but not identical, methods – Lower et al. (2020) used TEMPONEST rather than ENTERPRISE to perform Bayesian model selection and parameter estimation. In this section, we investigate the consistency of the estimates of timing noise parameters with respect to both timing baseline and methodology.

There is no obvious reason why the estimates of the timing noise PSD should differ significantly when ENTERPRISE is used instead of TEMPONEST, or when additional years of data are added. However, in the latter case, it is plausible that the PSD deviates from power-law behaviour as the frequency decreases, which would alter the inferred parameters when fitting to a pure power-law model as more data (and hence lower frequencies) are incorporated (Melatos & Link

Table 2. Estimates of the parameters determining χ_{RN} [equation (8)], as discussed in Section 4.1.

Parameter	This work ($f_{\text{low}}^{-\gamma}$)	This work [$f_{\text{low}}^{-(\beta-1)/2}$]	SC10	Lo + 20
$\log_{10} C$	$-2.6^{+2.5}_{-2.2}$	$-5.0^{+2.0}_{-1.6}$	2.0 ± 0.4	$3.7^{+2.4}_{-2.7}$
a	$-0.85^{+0.38}_{-0.35}$	$-0.49^{+0.31}_{-0.32}$	-0.9 ± 0.2	$-0.84^{+0.47}_{-0.49}$
b	0.56 ± 0.16	0.47 ± 0.13	1.00 ± 0.05	$0.97^{+0.16}_{-0.19}$
γ	1.2 ± 0.8	–	1.9 ± 0.2	1.0 ± 1.2
ϵ	0.73 ± 0.11	0.65 ± 0.1	0.69 ± 0.04	$0.64^{+0.11}_{-0.16}$

Notes. In this work, we consider two cases, one with $\chi_{\text{RN}} \propto f_{\text{low}}^{-\gamma}$ (second column), where γ is a free parameter to be inferred, and one with $\chi_{\text{RN}} \propto f_{\text{low}}^{-(\beta-1)/2}$ (third column). The fourth and fifth columns present the equivalent results from Shannon & Cordes (2010) and Lower et al. (2020), respectively. Error bars are 95 per cent credible intervals for the quoted results from this work and from Lower et al. (2020), and 2σ confidence intervals for the results of Shannon & Cordes (2010).




Figure 4. Distributions of the differences in recovered timing noise parameters A_{red} and β for various combinations of methodology and data set, as discussed in Section 4.2. The stars in each panel indicate the locations of the means of the joint distributions of differences – see Table 3 for the numerical values of these means. The means do not differ significantly from zero in any of the three cases (see Table 3), showing no evidence for a systematic offset in either $\log_{10} A_{\text{red}}$ or β .

Table 3. Mean differences in recovered timing noise parameters, $\langle \Delta \log_{10} A_{\text{red}} \rangle$ and $\langle \Delta \beta \rangle$, for three pairings of data set and method. The data set is either the combined UTMOST-EW and UTMOST-NS data set (EW + NS) or the UTMOST-EW data alone (EW). The method is either ENTERPRISE (ENT) or TEMPONEST (TEM). The (EW, TEM) timing noise parameters are taken from the analysis of Lower et al. (2020). The complete distributions of timing noise parameter differences are shown in Fig. 4. The quoted uncertainty on each mean difference is the standard error of the mean.

Data sets and methods	$\langle \Delta \log_{10} A_{\text{red}} \rangle$	$\langle \Delta \beta \rangle$
(EW + NS, ENT) – (EW, TEM)	-0.07 ± 0.04	-0.21 ± 0.15
(EW + NS, ENT) – (EW, ENT)	-0.02 ± 0.05	-0.24 ± 0.12
(EW, ENT) – (EW, TEM)	0.06 ± 0.04	-0.04 ± 0.16

2014; Caballero et al. 2016; Parthasarathy et al. 2019; Goncharov, Zhu & Thrane 2020; Antonelli, Basu & Haskell 2023).

The top panel of Fig. 4 shows the difference in $\log_{10} A_{\text{red}}$ and β between the estimates of this work, using ENTERPRISE and the combined EW + NS data set, and those of Lower et al. (2020), using TEMPONEST and the UTMOST-EW data alone. In this comparison, we include the pulsars which are common to both analyses and for which the timing noise is significant, according to the criterion $\sigma_{\text{RN}} > 3\sigma_{\text{ToA}}$ [see equation (7)]. For the purposes of these comparisons we do not use the TNLONGF2 model. We note that some of the spectral indices reported by Lower et al. (2020) are as large as 20. Here, we restrict our attention to pulsars from the Lower et al. (2020) sample which have $2 < \beta < 10$, matching the prior used on β in this work.

For notational convenience, we write the differences between the two sets of parameter estimates as $\Delta \log_{10} A_{\text{red}}$ and $\Delta \beta$ where it is clear from context which two sets of parameter estimates we are differencing. For example, in the comparison between the UTMOST-EW + NS ENTERPRISE analyses and the UTMOST-EW-only TEMPONEST analyses by Lower et al. (2020) we take the differences

$$\Delta \log_{10} A_{\text{red}} = \log_{10} A_{\text{red}}^{\text{EW+NS, ENT}} - \log_{10} A_{\text{red}}^{\text{EW, TEM}}, \quad (10)$$

$$\Delta \beta = \beta^{\text{EW+NS, ENT}} - \beta^{\text{EW, TEM}}, \quad (11)$$

where we abbreviate the software used as ENT (ENTERPRISE) and TEM (TEMPONEST).

The mean of the joint distribution of $(\Delta \log_{10} A_{\text{red}}, \Delta \beta)$, indicated by the star in the top panel of Fig. 4, is $(-0.07 \pm 0.04, -0.21 \pm 0.15)$ where the uncertainties are the standard errors of the mean. While there appears to be some scatter between the results obtained in the two sets of analyses, there is no evidence for a systematic offset between them.

However, the estimates in the top panel of Fig. 4 were obtained using a different method and a different time span. In order to eliminate the possibility of competing effects from changes in the two variables, we re-analyse the UTMOST-EW data using ENTERPRISE. The middle panel shows the difference between the ENTERPRISE results for the combined EW + NS data set and the ENTERPRISE results for the EW data alone, and the bottom panel shows the difference between the ENTERPRISE results for the EW data and the Lower et al. (2020) results. In the case of EW + NS with ENTERPRISE compared to EW-only with ENTERPRISE, there is at most marginal evidence for a shift in β , with $\langle \Delta \beta \rangle = -0.24 \pm 0.12$. The EW-only ENTERPRISE and TEMPONEST measurements appear to be consistent, with $\langle \Delta \beta \rangle = -0.04 \pm 0.16$. There is no evidence in either case for a systematic shift in $\Delta \log_{10} A_{\text{red}}$. The results of all three pairings discussed are summarized in Table 3.

The results of this section are re-assuring. The apparent lack of systematic variation in A_{red} and β as the analysis method and time span are varied suggests that it makes good sense to model the PSD of the timing noise in these pulsars as a simple power law, at least over the frequencies probed by our data sets. By extension, these results suggest that we may regard A_{red} and β as meaningful parameters in their own right, not only when combined into a figure of merit like σ_{RM} . For example, in Sections 5 and 6, we use the value of β as an indicator of the kind of random walk which is appropriate to model a given pulsar’s behaviour, and A_{red} gives an estimate of the variance of this random walk. Were we to observe systematic shifts in one or both of A_{red} and β as a function of the analysis method or time span, these interpretations of the PSD parameters would be less robust.

5 SECOND FREQUENCY DERIVATIVES

Timing noise and the influence of the second frequency derivative are both low-frequency effects, in the sense that their contributions to timing residuals are concentrated in Fourier modes with frequency comparable to T_{span}^{-1} . This can make it difficult to disentangle the influence of the two contributions.

The secular value of $\ddot{\nu}$ is of particular physical interest because of its connection to the long-term spin-down of the pulsar. For many spin-down mechanisms, the expected frequency evolution has the form

$$\dot{\nu} = K \nu^{n_{\text{pl}}}, \quad (12)$$

where, n_{pl} is referred to as the ‘braking index’ and may be estimated (as long as timing noise is low enough; see below) by $n_{\text{pl}} \approx n$ assuming K is a constant, with

$$n = \frac{\nu \ddot{\nu}}{\dot{\nu}^2}. \quad (13)$$

The value of n_{pl} differs between physical mechanisms (Blandford & Romani 1988; Vargas & Melatos 2023; Abolmasov, Biryukov & Popov 2024), but is typically in the range of $1 \lesssim n_{\text{pl}} \lesssim 7$, e.g. spin-down due to electromagnetic dipole radiation has $n_{\text{pl}} = 3$, spin-down due to gravitational wave emission from a mass quadrupole has $n_{\text{pl}} = 5$, and spin-down due to gravitational wave emission from a mass current quadrupole has $n_{\text{pl}} = 7$ (Lyne & Graham-Smith 2012; Riles 2023).

If one fits for $\ddot{\nu}$ in a pulsar with significant timing noise, one is liable to find large values of $\ddot{\nu}$, in the sense that the implied braking index is orders of magnitude different from the ‘canonical’ range of $1 \lesssim n \lesssim 7$ and can take either sign. It is well understood that such values are unreliable as indicators of the long-term behaviour, as the residuals due to timing noise can be absorbed into the $\ddot{\nu}$ term (Hobbs et al. 2004; Chukwude & Chidi Odo 2016). However, with the advent of Bayesian techniques which simultaneously fit for the residual contributions of the timing noise and the $\ddot{\nu}$ term, a number of ‘anomalous’ braking indices with large magnitudes have been reported, with values ranging from -9.6×10^4 to 2.9×10^3 (Lower et al. 2020; Parthasarathy et al. 2020). It has been suggested that these anomalous braking indices are due to the effects of recovery from glitches which occurred before the start of the data set, and do not necessarily reflect the long-term spin-down behaviour (Johnston & Galloway 1999; Lower et al. 2020, 2021; Parthasarathy et al. 2020). Table 4 lists all pulsars in our sample with a significant $\ddot{\nu}$ value, where ‘significant’ here means that the 95 per cent confidence interval for $\ddot{\nu}$ excludes zero.

There are 39 pulsars satisfying this condition. The measured values of n span a wide range, from $-2.5_{-1.2}^{+1.0} \times 10^4$ (PSR J2053–7200) to $8.4_{-7.1}^{+7.4} \times 10^4$ (PSR 1605–5257).

Recently, it has been pointed out that despite the more sophisticated handling of the mixing between timing noise and $\ddot{\nu}$ in these Bayesian analyses, reported values of n must still be treated with caution in some cases (Keith & Nițu 2023; Vargas & Melatos 2023). In both of these studies it was shown that under certain conditions, the phenomenological timing noise model of equation (3), particularly with a low-frequency cut-off at $1/T_{\text{span}}$ (which is the default setting in both ENTERPRISE and TEMPONEST), can lead to erroneous characterization of the $\ddot{\nu}$ term. In the remainder of this section, we discuss the interaction between timing noise and the estimation of uncertainty on $\ddot{\nu}$ (Section 5.1), the prospect of measuring n_{pl} for pulsars in our sample (Section 5.2), and the effect of the adopted timing noise model on the detection of $\ddot{\nu}$ (Section 5.3) in the context of our data set.

5.1 Timing noise and estimated $\ddot{\nu}$ uncertainties

We can estimate the expected spread in $\ddot{\nu}$ measurements due to the influence of timing noise and compare this to the $\ddot{\nu}$ uncertainties returned by ENTERPRISE in the TNF2 and TNLONGF2 models (Vargas & Melatos 2023). Timing noise causes the observed $\ddot{\nu}$ to change from realization to realization, but only one realization is observed in an astronomical observation, and it is impossible to determine independently whether the observed realization is typical or rare with in the ensemble of physically permitted realizations.

In the first instance, we take the subset of pulsars for which the timing noise parameter estimation for both the TNF2 and TNLONGF2 models returns a spectral index satisfying $|\beta - 6| < 1$. This condition on β suggests that the timing noise may be approximately modelled by a random walk in $\dot{\nu}$, i.e.³

$$\frac{d\dot{\nu}}{dt} = \ddot{\nu}_0 + \xi(t), \quad (14)$$

where $\ddot{\nu}_0$ is an intrinsic constant second frequency derivative and $\xi(t)$ is a white noise term satisfying

$$\langle \xi(t) \rangle = 0, \quad (15)$$

$$\langle \xi(t)\xi(t') \rangle = \sigma_{\dot{\nu}}^2 \delta(t - t'). \quad (16)$$

This choice corresponds to a particular limit of the model used by Vargas & Melatos (2023). There, the stochastic driving term is in $d\dot{\nu}/dt$, and $\ddot{\nu}$ and $\dot{\nu}$ revert to their means on time-scales $\gamma_{\dot{\nu}}^{-1}$ and $\gamma_{\ddot{\nu}}^{-1}$, respectively [see equations (2)–(5) of the latter reference]. The model given by equations (14)–(16) here corresponds to the limit $\gamma_{\dot{\nu}}^{-1} \gg T_{\text{span}}$ and $\gamma_{\ddot{\nu}}^{-1} \lesssim T_{\text{cad}}$ where T_{cad} is the observing cadence; cf. equations (A15) and (A16) of Vargas & Melatos (2023).

The PSD of the phase, $P_{\phi}(f)$, may be estimated by the Wiener–Kinchin theorem, which connects the autocorrelation function with the PSD via a Fourier transform:

$$P_{\phi}(f) = 2\sigma_{\dot{\nu}}^2 (2\pi f)^{-6}. \quad (17)$$

We then wish to compare $P_{\phi}(f)$ to $P(f)$ inferred from ENTERPRISE which is parametrized as in equation (3), and hence obtain an

³For simplicity, we do not include a mean-reverting term in the equation of motion, on the assumption that the mean-reversion time-scale for $\dot{\nu}$ is significantly longer than the observing time span (Vargas & Melatos 2023).

Table 4. Pulsars for which we measure a non-zero $\ddot{\nu}$ with at least 95 per cent confidence.

PSR	$\ddot{\nu}$ (Hz s ⁻²)	n	$\log_{10} A_{\text{red}}$	β	Model
J0401–7608	$-9.5^{+5.2}_{-6.5} \times 10^{-26}$	$-6.5^{+3.5}_{-4.4} \times 10^3$	-11.7	6.66	TNF2
J0534+2200	$(1.16 \pm 0.01) \times 10^{-20}$	2.54 ± 0.03	-9.21	4.1	TNF2
J0630–2834	$(-6.5 \pm 0.7) \times 10^{-26}$	$(-2.4 \pm 0.3) \times 10^3$	-12.6	5.33	TNF2
J0758–1528	$2.7^{+3.1}_{-2.6} \times 10^{-26}$	$3.2^{+3.7}_{-3.1} \times 10^3$	-11.9	7.02	TNF2
J0835–4510	$(1.23 \pm 0.01) \times 10^{-21}$	57.4 ± 0.6	-8.47	7.44	TNF2
J0953+0755	$-6.3^{+3.0}_{-2.5} \times 10^{-26}$	$-1.9^{+0.9}_{-0.8} \times 10^4$	-10.8	3.09	TNF2
J0955–5304	$4.5^{+3.5}_{-3.4} \times 10^{-26}$	$2.3^{+1.8}_{-1.7} \times 10^3$	-11.7	6.88	TNF2
J1001–5507	$9.6^{+4.3}_{-4.1} \times 10^{-25}$	$(1.1 \pm 0.5) \times 10^3$	-9.25	4.39	TNF2
J1017–5621	$8.3^{+2.2}_{-2.0} \times 10^{-26}$	$(1.1 \pm 0.3) \times 10^3$	-12.5	8.02	TNF2
J1048–5832	$(1.1 \pm 0.3) \times 10^{-22}$	$23.2^{+5.4}_{-5.8}$	-8.7	4.55	TNLF2
J1136–5525	$-2.1^{+1.2}_{-1.1} \times 10^{-24}$	$-1.5^{+0.9}_{-0.8} \times 10^3$	-10	7.16	TNF2
J1141–6545	$2.2^{+1.1}_{-1.2} \times 10^{-25}$	$7.1^{+3.7}_{-3.9} \times 10^2$	-11.4	6.93	TNF2
J1202–5820	$(-2.6 \pm 2.5) \times 10^{-25}$	$-5.2^{+5.1}_{-5.0} \times 10^3$	-10.6	5.98	TNF2
J1326–6408	$-5.2^{+2.7}_{-2.6} \times 10^{-26}$	$-2.7^{+1.4}_{-1.3} \times 10^3$	-12	7.09	TNF2
J1327–6222	$7.8^{+7.0}_{-7.2} \times 10^{-25}$	$(3.3 \pm 3.0) \times 10^2$	-9.12	3.51	TNLF2
J1328–4357	$(1.9 \pm 1.7) \times 10^{-25}$	$(3.1 \pm 2.8) \times 10^3$	-10.5	5.39	TNF2
J1507–4352	$-1.6^{+1.6}_{-1.8} \times 10^{-25}$	$-1.5^{+1.5}_{-1.7} \times 10^3$	-10.1	3.44	TNF2
J1522–5829	$(-4.1 \pm 0.4) \times 10^{-25}$	$-6.0^{+0.6}_{-0.5} \times 10^3$	-10.9	3.69	TNF2
J1553–5456	$7.1^{+6.5}_{-6.1} \times 10^{-25}$	$3.7^{+3.3}_{-3.2} \times 10^3$	-10.4	7.61	TNF2
J1605–5257	$1.9^{+1.9}_{-1.7} \times 10^{-26}$	$8.4^{+8.2}_{-7.2} \times 10^4$	-12.2	6.64	TNF2
J1644–4559	$8.9^{+5.6}_{-5.2} \times 10^{-25}$	$2.1^{+1.3}_{-1.2} \times 10^2$	-9.95	5.7	TNF2
J1645–0317	$(-1.3 \pm 1.2) \times 10^{-24}$	$-2.5^{+2.2}_{-2.3} \times 10^4$	-10	8.46	TNF2
J1651–4246	$-6.7^{+3.1}_{-2.9} \times 10^{-25}$	$(-1.8 \pm 0.8) \times 10^4$	-10.4	7.03	TNF2
J1651–5222	$4.1^{+4.0}_{-3.5} \times 10^{-26}$	$3.2^{+3.1}_{-2.7} \times 10^3$	-11.8	6.54	TNF2
J1703–3241	$8.2^{+2.5}_{-3.1} \times 10^{-27}$	$3.4^{+1.0}_{-1.3} \times 10^4$	-12.8	5.11	TNF2
J1709–1640	$4.7^{+2.1}_{-1.9} \times 10^{-25}$	$3.3^{+1.5}_{-1.4} \times 10^3$	-10.2	4.99	TNF2
J1709–4429	$(2.1 \pm 0.2) \times 10^{-22}$	$26.1^{+2.8}_{-2.4}$	-9.62	6.49	TNF2
J1731–4744	$(1.3 \pm 0.2) \times 10^{-24}$	27.7 ± 5.1	-9.79	3.61	TNF2
J1740–3015	$8.2^{+2.8}_{-2.4} \times 10^{-24}$	$8.4^{+2.9}_{-2.4}$	-9	5.27	TNF2
J1741–3927	$1.0^{+0.2}_{-0.1} \times 10^{-24}$	$(5.0 \pm 0.7) \times 10^4$	-11.5	9.07	TNF2
J1743–3150	$2.9^{+1.3}_{-2.3} \times 10^{-26}$	$27.7^{+12.6}_{-22.1}$	-12.2	6	TNF2
J1745–3040	$(1.8 \pm 0.4) \times 10^{-25}$	$77.4^{+16.9}_{-17.9}$	-11.1	4.63	TNF2
J1752–2806	$(-1.8 \pm 1.2) \times 10^{-25}$	$-4.8^{+3.3}_{-3.1} \times 10^2$	-9.57	2.83	TNF2
J1757–2421	$7.0^{+3.5}_{-4.4} \times 10^{-25}$	$53.7^{+26.5}_{-33.7}$	-10.5	5.02	TNF2
J1823–3106	$5.8^{+1.4}_{-1.5} \times 10^{-25}$	$(1.5 \pm 0.4) \times 10^3$	-10.4	4.11	TNF2
J1824–1945	$(1.2 \pm 0.8) \times 10^{-24}$	$(3.1 \pm 2.0) \times 10^2$	-10.1	5.49	TNF2
J1847–0402	$(1.4 \pm 0.2) \times 10^{-25}$	$11.0^{+1.4}_{-2.0}$	-11.1	3.93	TNF2
J1848–0123	$2.5^{+0.8}_{-0.7} \times 10^{-25}$	$2.6^{+0.8}_{-0.7} \times 10^3$	-10.8	5.35	TNF2
J2053–7200	$(-2.5 \pm 1.1) \times 10^{-26}$	$(2.5 \pm 1.1) \times 10^4$	-12.8	5.22	TNF2

Notes. The measured value of $\ddot{\nu}$ is given in the second column, while the implied value of n [equation (13)] appears in the third. The estimated properties of the timing noise PSD, A_{red} and β , are given in the fourth and fifth columns. The model chosen according to the criteria defined in Section 3.2 is listed in the final column.

expression for $\sigma_{\ddot{\nu}}$ in terms of A_{red} and β . In order to render $P_{\phi}(f)$ and $P(f)$ comparable, there are two factors which must be taken into account. First, $P_{\phi}(f)$ is the PSD of the residuals in units of cycles, while $P(f)$ is the PSD of the residuals in units of seconds. Thus, $P(f)$ should be multiplied by a factor of v^2 in order to match $P_{\phi}(f)$. Second, $P_{\phi}(f)$ has a spectral index of exactly -6 , whereas $P(f)$ generally does not. For the purposes of this comparison we therefore replace $P(f)$ with a new PSD $P^{\beta=6}(f)$ parametrized in the same way as $P(f)$, but with β fixed to 6 and A_{red} adjusted so that the values of σ_{RN} [equation (7)] for $P(f)$ and $P^{\beta=6}(f)$ coincide. This

adjusted A_{red} value is denoted by $A_{\text{red}}^{\beta=6}$, and is given by

$$A_{\text{red}}^{\beta=6} = \sqrt{\frac{5}{\beta-1}} A_{\text{red}} \left(\frac{f_{\text{low}}}{f_{\text{yr}}} \right)^{-(\beta-6)/2}. \quad (18)$$

Equating $P_{\phi}(f)$ and $v^2 P^{\beta=6}(f)$ and isolating $\sigma_{\ddot{\nu}}$ gives

$$\sigma_{\ddot{\nu}} = \sqrt{\frac{f_{\text{yr}}^3 v^2 (2\pi)^6}{24\pi^2}} A_{\text{red}}^{\beta=6}. \quad (19)$$

Given $\sigma_{\dot{\nu}}$, we wish to estimate the spread of observed $\dot{\nu}$ values across different realizations of the noise process. We denote this spread by $\delta\dot{\nu}$. Over an observation of length T , $\delta\dot{\nu}$ can be estimated by calculating the dispersion in $\dot{\nu}$ due to the random walk, given by $\sigma_{\dot{\nu}}T^{1/2}$, and dividing this value by T , i.e.

$$\delta\dot{\nu} \sim \sqrt{\frac{f_{\text{yr}}^3 \nu^2 (2\pi)^6}{24\pi^2}} A_{\text{red}}^{\beta=6} T^{-1/2}. \quad (20)$$

If the $\dot{\nu}$ which we infer is to represent the long-term spin-down trend, $\delta\dot{\nu}$ ought to give an approximate lower bound on the reported uncertainty on $\dot{\nu}$ for pulsars with timing noise following equations (14)–(16). Any long-term secular $\dot{\nu}$ smaller than this value is liable to be washed out by the influence of the timing noise.

We compare the value of $\delta\dot{\nu}$ to $\Delta\dot{\nu}$, the estimated uncertainty on $\dot{\nu}$ returned by ENTERPRISE for all pulsars with $\sigma_{\text{RN}} > 3\sigma_{\text{ToA}}$. As in Section 4.1, this restriction excludes pulsars with low-amplitude timing noise whose PSDs are poorly measured. For the TNF2 case, the median value of $\delta\dot{\nu}/\Delta\dot{\nu}$ is 2.59, while in the TNLONGF2 case it is 0.91. The TNF2 model tends to underestimate $\Delta\dot{\nu}$ relative to the expected spread in $\dot{\nu}$ due to timing noise, $\delta\dot{\nu}$, while the $\dot{\nu}$ uncertainty estimated by the TNLONGF2 model better reflects this variance.

We repeat this analysis for pulsars with $|\beta - 4| < 1$, suggesting a Langevin equation of the form

$$\frac{d\nu}{dt} = \dot{\nu}_0 + \xi(t), \quad (21)$$

with the strength of the $\xi(t)$ term now parametrized by σ_{ν} rather than $\sigma_{\dot{\nu}}$. By analogous arguments to the $|\beta - 6| < 1$ case, we find that the expected spread in observed $\dot{\nu}$ is given by (details are given in Appendix B)

$$\delta\dot{\nu} \sim 2\sqrt{\frac{f_{\text{yr}} \nu^2 (2\pi)^4}{24\pi^2}} A_{\text{red}}^{\beta=4} T^{-3/2}. \quad (22)$$

Again comparing $\delta\dot{\nu}$ with $\Delta\dot{\nu}$, we find that the mean $\delta\dot{\nu}/\Delta\dot{\nu}$ in our sample is 0.50 and 0.40 for TNF2 and TNLONGF2 respectively – i.e. both timing noise models appear to overestimate the uncertainty in $\dot{\nu}$ relative to the variance due to timing noise, but the degree of the overestimation is less sensitive to f_{low} for $|\beta - 4| < 1$ than for $|\beta - 6| < 1$. Fig. 5 shows histograms of $\log_{10}(\delta\dot{\nu}/\Delta\dot{\nu})$ for $|\beta - 6| < 1$ (top) and $|\beta - 4| < 1$ (bottom). Each panel shows two histograms of $\log_{10}(\delta\dot{\nu}/\Delta\dot{\nu})$, one for TNF2 and the other for TNLONGF2. For $|\beta - 6| < 1$, there is a clear separation between the histograms for the two timing noise models, with the TNF2 model tending to have larger values of $\delta\dot{\nu}/\Delta\dot{\nu}$, i.e. the spread due to timing noise exceeds the reported uncertainty, as noted above. For $|\beta - 4| < 1$, both models produce similar histograms and typically have $\delta\dot{\nu}/\Delta\dot{\nu} < 1$.

These results broadly agree with the conclusions of Keith & Nițu (2023) – pulsars with $\beta \sim 6$ suffer more from mis-estimation of $\dot{\nu}$ uncertainty than those with $\beta \sim 4$, especially if a timing noise model with a low-frequency cut-off at $1/T_{\text{span}}$ is used.

5.2 Measuring n_{pl}

The values of $\delta\dot{\nu}$ obtained in Section 5.1 can also be cast in terms of the measured braking index. We define

$$\delta n = \frac{\nu \delta\dot{\nu}}{\dot{\nu}^2}, \quad (23)$$

as the expected variation in measured n due to the influence of timing noise. This raises an interesting question: if every pulsar in our data set $1 \lesssim n_{\text{pl}} \lesssim 7$, for how many pulsars might we expect to measure n_{pl} , given our current data set? And if timing noise currently




Figure 5. Histograms comparing the expected variation in the measured value of $\dot{\nu}$ due to timing noise, $\delta\dot{\nu}$, against the uncertainty reported by ENTERPRISE, $\Delta\dot{\nu}$. The results are divided into two cases depending on the estimated spectral index of the PSD, with $|\beta - 6| < 1$ in the top panel and $|\beta - 4| < 1$ in the bottom panel. Two histograms are shown in each panel, corresponding to the TNF2 ($f_{\text{low}} = 1/T_{\text{obs}}$) and TNLONGF2 ($f_{\text{low}} = 1/2T_{\text{obs}}$) timing noise models.

confounds a measurement of n_{pl} , for how long do we need to observe to overcome this effect? Vargas & Melatos (2023) derived a condition for when we measure $|n| \gg n_{\text{pl}}$ in terms of the strength of the noise process, the pulsar spin frequency and its first derivative, and the observing time span; see equation (14) in Vargas & Melatos (2023). Their condition follows from $\delta n > 1$. Here, we write the equivalent condition on the noise strength for $|\beta - 6| < 1$ [which is a condition on $\sigma_{\dot{\nu}}$, see Section 5.1], viz.

$$\sigma_{\dot{\nu}} > 10^{-20} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-1} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{1/2} \text{ s}^{-5/2}, \quad (24)$$

and $|\beta - 4| < 1$ [which is a condition on σ_{ν} , see Appendix B], viz.

$$\sigma_{\nu} > 10^{-12} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-1} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{3/2} \text{ s}^{-3/2}. \quad (25)$$

These can be rephrased as conditions on $A_{\text{red}}^{\beta=6}$, viz.

$$A_{\text{red}}^{\beta=6} > 1.1 \times 10^{-10} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-2} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{1/2}, \quad (26)$$

Table 5. Pulsars with $\delta n < 1$, where δn is the expected variation due to timing noise in the measured value of n , as discussed in Section 5.1. The values of $\log_{10} A_{\text{red}}$, β , $\dot{\nu}$, and n are taken from the TNLONGF2 model, and the values of $\delta\dot{\nu}$ and δn are calculated according to equations (22) and (23). The quoted uncertainties on $\dot{\nu}$ and n are 95 per cent credible intervals.

PSR	$\log_{10} A_{\text{red}}$	β	T_{span} (yr)	$\dot{\nu}$ (Hz s^{-2})	$\delta\dot{\nu}$ (Hz s^{-2})	n	δn
J0534+2200	-9.19	4.1	1.9	$(1.16 \pm 0.02) \times 10^{-20}$	3.7×10^{-23}	$2.54^{+0.04}_{-0.03}$	0.0082
J1731-4744	-9.96	3.68	8.2	$(1.2 \pm 0.3) \times 10^{-24}$	2.9×10^{-26}	26.7 ± 6.7	0.63




Figure 6. Distribution of T_{span} required to obtain $\delta n < 1$ [equation (23)] using the TNLONGF2 model, for the cases where the spectral index of the timing noise PSD satisfies $|\beta - 6| < 1$ and $|\beta - 4| < 1$. The median required T_{span} is 7.9×10^7 yr for $|\beta - 6| < 1$, much longer than the 2.8×10^2 yr for $|\beta - 4| < 1$.

and $A_{\text{red}}^{\beta=4}$, viz.

$$A_{\text{red}}^{\beta=4} > 1.1 \times 10^{-9} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-2} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{3/2}. \quad (27)$$

With equations (24)–(27) in hand, we ask: for which pulsars are the relevant inequalities not satisfied? That is, for which pulsars do we expect the variation in the measured value of n due to the influence of timing noise, δn [see equation (23)], to be less than unity? Table 5 lists the two pulsars with $\delta n < 1$ over the time span covered by the combined EW + NS data set, using the parameters returned by the TNLONGF2 model.

For all other pulsars, we do not expect to be able to resolve $\dot{\nu}$ due to ‘canonical’ secular spin-down with the UTMOST-NS + EW data. However, δn is a decreasing function of T for both $|\beta - 4| < 1$ and $|\beta - 6| < 1$, so we expect that with a sufficiently long timing baseline, n_{pl} can be measured [again neglecting for simplicity any mean reversion (Vargas & Melatos 2023)]. Fig. 6 shows histograms of the required T_{span} to obtain $\delta n < 1$ for the cases $|\beta - 6| < 1$ (blue) and $|\beta - 4| < 1$ (orange).

Pulsars with $|\beta - 6| < 1$ require significantly longer T_{span} on average – the median required T_{span} is 7.9×10^7 yr, compared to 2.8×10^2 yr for $|\beta - 4| < 1$.

5.3 Influence of frequency cut-off in timing noise models on the detection of $\dot{\nu}$

As discussed at the beginning of Section 5, a total of 39 pulsars in our sample have posterior $\dot{\nu}$ distributions which exclude zero at 95 per cent confidence. The results of Section 5.1 indicate that TNF2 analyses can underestimate the uncertainty in $\dot{\nu}$ significantly. This

Table 6. Non-zero $\dot{\nu}$ measurements in our sample when the TNF2 model is removed from consideration (see Section 5.3). As in Table 4, all pulsars whose 95 per cent confidence interval on $\dot{\nu}$ excludes zero are included.

PSR	$\dot{\nu}$ (Hz s^{-2})	n	$\log_{10} A_{\text{red}}$	β
J0534+2200	$(1.16 \pm 0.02) \times 10^{-20}$	$2.54^{+0.04}_{-0.03}$	-9.21	4.1
J0630-2834	$-6.5^{+2.1}_{-1.8} \times 10^{-26}$	$-2.4^{+0.8}_{-0.7} \times 10^3$	-12.5	6.15
J0953+0755	$-6.2^{+3.5}_{-2.9} \times 10^{-26}$	$-1.9^{+1.1}_{-0.9} \times 10^4$	-10.8	3.15
J1001-5507	$9.6^{+5.2}_{-5.0} \times 10^{-25}$	$(1.1 \pm 0.6) \times 10^3$	-9.25	4.39
J1017-5621	$7.5^{+8.0}_{-6.7} \times 10^{-26}$	$1.0^{+1.0}_{-0.9} \times 10^3$	-13	8.81
J1048-5832	$(1.1 \pm 0.3) \times 10^{-22}$	$23.2^{+5.4}_{-5.8}$	-8.7	4.55
J1327-6222	$7.8^{+7.0}_{-7.2} \times 10^{-25}$	$(3.3 \pm 3.0) \times 10^2$	-9.12	3.51
J1507-4352	$-1.6^{+1.4}_{-1.8} \times 10^{-25}$	$-1.5^{+1.3}_{-1.7} \times 10^3$	-10.1	3.4
J1522-5829	$-4.1^{+0.7}_{-0.6} \times 10^{-25}$	$-5.9^{+1.0}_{-0.8} \times 10^3$	-10.9	3.74
J1709-1640	$4.6^{+4.3}_{-4.0} \times 10^{-25}$	$3.2^{+3.0}_{-2.8} \times 10^3$	-10.2	5
J1709-4429	$2.0^{+0.5}_{-0.6} \times 10^{-22}$	$25.1^{+6.0}_{-7.5}$	-9.6	6.19
J1731-4744	$(1.2 \pm 0.3) \times 10^{-24}$	26.7 ± 6.7	-9.78	3.68
J1745-3040	$1.8^{+0.6}_{-0.7} \times 10^{-25}$	$77.0^{+26.8}_{-29.8}$	-11.1	4.62
J1752-2806	$-1.7^{+1.1}_{-1.0} \times 10^{-25}$	$-4.5^{+3.0}_{-2.8} \times 10^2$	-9.58	2.84
J1823-3106	$5.8^{+1.7}_{-1.9} \times 10^{-25}$	$(1.5 \pm 0.5) \times 10^3$	-10.4	4.18
J1847-0402	$(1.4 \pm 0.4) \times 10^{-25}$	$10.8^{+3.2}_{-3.4}$	-11.1	4.28
J1848-0123	$2.5^{+2.0}_{-1.3} \times 10^{-25}$	$2.7^{+2.1}_{-1.3} \times 10^3$	-10.8	5.32

motivates us to re-examine the measured $\dot{\nu}$ values in our sample, but with the restriction that only the TNLONGF2 model be considered. Table 6 shows the list of pulsars for which we recover a $\dot{\nu}$ posterior which excludes zero at 95 per cent confidence in this case. There are 17 in total – excluding the TNF2 model removes approximately half the $\dot{\nu}$ detections in our sample.

We may consider one further restriction. We construct an alternative model, TNLONG, which matches TNLONGF2 except that $\dot{\nu}$ is fixed at zero. If we further restrict our attention to pulsars for which TNLONGF2 is favoured over TNLONG by a log Bayes factor of 5 or more, there is just one pulsar satisfying these criteria: J0534+2200 (the Crab pulsar). The measured n of 2.54 ± 0.03 agrees with the interglitch value of 2.519(2) reported by Lyne et al. (2015).

Of course, there is no reason that the characteristics of the timing noise, as a physical process, should be sensitive to the observing span. As such, one might regard the TNF2 model as imposing an artificial restriction on the character of the timing noise, by enforcing a cut-off in the PSD at $1/T_{\text{span}}$. On the other hand, TNLONG and TNLONGF2 are necessarily more complex than TNF2, having additional sinusoidal components, and so are disfavoured by Occam’s razor (although this is incorporated naturally into the Bayes factor calculations). The proper interpretation of the $\dot{\nu}$ measurements in this section remains unclear – some may be spurious, while others may represent real physical processes other than the long-term braking, e.g. glitch recoveries (Hobbs et al. 2010; Parthasarathy et al. 2020; Lower et al. 2021; Liu et al. 2024b). We defer further investigation

into this issue to future work, and do not interpret the $\dot{\nu}$ values in Tables 4 and 6 further.

6 GLITCHES

A significant motivation for the UTMOST-NS pulsar timing programme is to monitor the pulsars in the programme for glitching activity. With approximately 50 pulsars per day being observed, manually inspecting every new observation for a possible glitch and responding in a timely manner is not possible, and an automated solution to ‘online’ glitch detection is implemented using the hidden Markov model (HMM)-based glitch detector⁴ introduced by Melatos et al. (2020). In addition, a subsequent ‘offline’ glitch search combining data from UTMOST-EW and UTMOST-NS is carried out, again using the HMM glitch detector. After an introduction to the HMM glitch detector, these two glitch searches, and their results, are described in the rest of this section.

6.1 Glitch detection with an HMM

We begin with a brief overview of how the HMM-based glitch detector operates. The reader is referred to Melatos et al. (2020) for a full description of the method, and Lower et al. (2021), Dunn et al. (2022a, 2023a) for previous examples of the use of this method to search successfully for glitches.

HMMs provide a convenient way to model systems in which the internal state of the system evolves in a Markovian (i.e. memoryless) fashion, and in which we are unable to observe the internal state directly, instead making proxy measurements which are connected probabilistically to the true internal states (i.e. the internal state is ‘hidden’) (Rabiner 1989). In this paper, the internal state of the system (pulsar) is taken to be the tuple of the spin frequency and frequency derivative $q_i = (\nu_i, \dot{\nu}_i)$. The hidden state is indexed with the discrete variable i , as the space of possible hidden states is discretized on a bounded grid of ν and $\dot{\nu}$ values referred to as the ‘domain of interest’ (DOI) (see Section 6.2.1). Time is discretized into a set of time-steps $\{t_1, t_2, \dots, t_{N_T}\}$, where each time-step corresponds to a gap between consecutive ToAs. We assume that the hidden state of the pulsar is fixed during each gap, but may change at the boundary between gaps according to some prescribed Markovian model, discussed in more detail in Section 6.2.2, so that we have a sequence of hidden states $\{q(t_1), q(t_2), \dots, q(t_{N_T})\}$.

The connection between the observed ToAs and the $(\nu, \dot{\nu})$ states is provided by the emission likelihood which is the probability of observing a particular gap between consecutive ToAs, given the pulsar occupies some hidden state q_i during this gap. Explicitly, the emission likelihood takes the form of a von Mises distribution,

$$\Pr(z | q_i) = \frac{\exp[\kappa \cos(2\pi \Phi)]}{2\pi I_0(\kappa)}, \quad (28)$$

where Φ denotes the number of turns accumulated during a gap of length z given the pulsar is in the hidden state q_i , and κ is a parameter known as the ‘concentration’, which parametrizes the uncertainty due to both ToA measurement uncertainty and the discretization of the space of hidden states. The reader is referred to section 3 of Dunn et al. (2022a) for further discussion of the forms of Φ and κ . The emission likelihood peaks, where the fractional part of Φ is zero, i.e. choices of ν and $\dot{\nu}$ for which an integer number of turns elapse during the ToA gap.

⁴github.com/ldunn/glitch_hmm

In this formalism, it is straightforward to compute the model evidence $\Pr(D | \mathcal{M})$, where D denotes the observed data and \mathcal{M} is a chosen model, which among other things incorporates assumptions about the Markovian transitions from one time-step to the next. This leads to a method of glitch detection based on Bayesian model selection, wherein we compare a model in which the transition dynamics include only an assumed Markov process representing the timing noise, denoted by \mathcal{M}_0 , against a series of models which include both the Markov process and the possibility of a discrete jump in frequency and frequency derivative (i.e. a glitch) during the k th ToA gap, denoted by $\mathcal{M}_1(k)$. We compute the Bayes factors

$$K_1(k) = \frac{\Pr[D | \mathcal{M}_1(k)]}{\Pr[D | \mathcal{M}_0]} \quad (29)$$

for all k . If the maximum value of $K_1(k)$, denoted by $K_1^* = K_1(k^*)$, exceeds a pre-determined threshold, here taken to be $10^{1/2}$, we flag a glitch candidate. Multiple glitches in one data set can be handled via a greedy procedure; see section 4.2 of Melatos et al. (2020) for details.

We measure the properties of a detected glitch using the forward–backward algorithm (Rabiner 1989) to efficiently compute the posterior distribution, i.e. $\gamma_{q_i}(t_n) = \Pr[q(t_n) = q_i | D]$. This distribution can be marginalized over either $\dot{\nu}$ or ν to obtain the posterior distribution on ν or $\dot{\nu}$ alone, respectively. Hence we can compute trajectories in ν or $\dot{\nu}$ by taking the sequence of a posteriori most probable hidden states, and from these trajectories we may compute e.g. the size of the frequency jump associated with a glitch.

As an illustrative example, Fig. 7 shows the results of the HMM analysis of a new small glitch in PSR J1902+0615 (see Section 6.5.13).

The top panel shows the log Bayes factors obtained as part of the model selection procedure. The blue points are the values of $K_1(k)$ [see equation (29)], which peak at $k^* = 57$ (i.e. the 57th ToA gap) with a maximum value of $\ln K_1^* = 15.9$. The orange points show the results of the second iteration of the greedy algorithm, where we compare the two-glitch model $\mathcal{M}_2(k^*, k)$ against the one-glitch model $\mathcal{M}_1(k^*)$ for all $k \neq k^*$. The largest log Bayes factor $K_2(k^*, k)$ obtained in this second step is 0.32, below the threshold $\ln 10^{1/2} = 1.15$. Thus, the model selection terminates and we take $\mathcal{M}_1(57)$ as our preferred model. The middle and bottom panels of Fig. 7 show heat maps of the posterior distributions of the hidden states ν and $\dot{\nu}$. The maximum a posteriori values are overlaid as red tracks, and the location of the glitch is indicated by the vertical dashed line. The jump in ν is clear, occurring around MJD 58 590 with size $(1.1 \pm 0.1) \times 10^{-9}$ Hz. No change in $\dot{\nu}$ is apparent at the time of the glitch.

6.2 HMM parameters

Completely specifying the HMM requires choosing a number of parameters controlling what set of $(\nu, \dot{\nu})$ pairs are allowed (i.e. the DOI), and how the hidden states evolve over time due to timing noise. In this section, we describe some general considerations influencing both choices. The specifics of how the parameters are set for the online and offline glitch searches are discussed in Section 6.3 and Appendix C, respectively.

6.2.1 Domain of interest

The DOI is the set of hidden states available to the HMM in a given analysis. The hidden states are tuples of frequency and frequency derivative $(\nu, \dot{\nu})$. For practical reasons, we define the hidden states




Figure 7. Illustrative example of a glitch analysis for PSR J1902+0615 using the HMM pipeline (Melatos et al. 2020), discussed in Section 6.1. The top panel shows the results of the model selection procedure, with log Bayes factors $K_1(k)$ (comparing one-glitch models against the no-glitch model; k indexes the ToA gap containing a glitch) and $K_2(57, k)$ (comparing two-glitch models against the best one-glitch model) plotted versus k (matched to MJD). The peak one-glitch log Bayes factor is 15.9, well above the threshold of 1.15, while the peak two-glitch log Bayes factor is 0.32. The middle and bottom panels show heat maps of the posterior distribution of the hidden states ν and $\dot{\nu}$, with the maximum a posteriori tracks overlaid. The estimated size of the glitch based on the ν posterior distribution is $(1.1 \pm 0.1) \times 10^{-9}$ Hz. There is no discernible jump in $\dot{\nu}$. The vertical dashed lines indicate the location of the glitch. The values of γ_ν and $\gamma_{\dot{\nu}}$ are normalized so they sum to unity at each time-step.

nu ($\dot{\nu}$) to be the difference between the spin frequency (frequency derivative) and its approximate value as determined by an initial least-squares fit, denoted by ν_{LS} ($\dot{\nu}_{\text{LS}}$). The least-squares values are defined at a reference epoch T_0 . The phase accumulated during a gap of length z ending at time t_n is thus given by

$$\Phi(z, t_n) = [\nu_{\text{LS}} + \nu + \dot{\nu}_{\text{LS}}(t_n - T_0)]z - \frac{1}{2}(\dot{\nu}_{\text{LS}} + \dot{\nu})z^2, \quad (30)$$

where the second term includes a minus sign because we are performing a backwards Taylor expansion (see section 3.3 of Melatos et al. 2020). A secular second frequency derivative $\ddot{\nu}_{\text{LS}}$ can also be included, if necessary. Note that because the spin-down behaviour is included via the secular terms, we typically take the DOI to be symmetric about $(\nu, \dot{\nu}) = (0, 0)$.

In all cases in this work, the DOI is laid out as a uniform grid in $(\nu, \dot{\nu})$. Specifying the DOI thus amounts to specifying the bounds and grid spacing in each parameter. This varies on a case-by-case basis and is discussed in more detail in Section 6.3 for the online detection case, and Appendix C for the offline detection case. Here, we briefly introduce the notation used to describe the DOI for ν ; for $\dot{\nu}$ the notation is identical but replacing ν with $\dot{\nu}$. The bounds of the DOI in ν are denoted by $[\nu_-, \nu_+]$. As mentioned, the DOI is typically symmetric about $(0, 0)$, whereupon we have $\nu_- = -\nu_+$ and write the bounds as $[-\nu_\pm, \nu_\pm]$. The grid spacing is denoted by ϵ_ν , i.e. ν takes on values $\{\nu_-, \nu_- + \epsilon_\nu, \nu_- + 2\epsilon_\nu, \dots, \nu_+\}$.

6.2.2 Timing noise models

In principle, the HMM specified in Section 6.1 can be used with various models for the underlying hidden state evolution, which amounts to specifying the model for the timing noise in the pulsar. This matches the situation in traditional, non-HMM glitch searches. Many past analyses assume that the stochastic evolution is driven by a white noise term in the second frequency derivative, i.e. a random walk in $\dot{\nu}$. This model is identical to the one specified by equations (14)–(16). We call it RWF1. The form of the transition matrix $A_{q_i q_j} = \text{Pr}[q(t_{n+1}) = q_j | q(t_n) = q_i, o(t_n) = z]$ is a multivariate Gaussian with the covariances

$$\text{cov}(\nu, \nu) = \sigma_\nu^2 z^3 / 3, \quad (31)$$

$$\text{cov}(\dot{\nu}, \nu) = \sigma_\nu^2 z^2 / 2, \quad (32)$$

$$\text{cov}(\dot{\nu}, \dot{\nu}) = \sigma_\nu^2 z. \quad (33)$$

The reader is referred to section 3.4 of Melatos et al. (2020) for the full expression for $A_{q_i q_j}$.

As discussed in Section 5.1, the RWF1 model gives rise to phase residuals with a PSD which has a spectral index of -6 . While many pulsars do exhibit residuals of this sort, a significant fraction also exhibit shallower PSDs. This motivates us to introduce a second model for timing noise in the HMM, this time driven by a white noise term in the frequency derivative, which corresponds to a random walk in ν and which gives rise to residuals with a spectral index of -4 . We call this model RWF0, and its Langevin equation is

$$\frac{d\nu}{dt} = \dot{\nu}_0 + \xi(t), \quad (34)$$

where $\xi(t)$ again satisfies equations (15) and (16), except that the parameter determining the noise amplitude is labelled σ_ν instead of $\sigma_{\dot{\nu}}$. Writing down the covariances of the Gaussian describing the transition matrix in this case requires some care – the autocorrelation of $\dot{\nu}$ diverges, with $\langle \dot{\nu}(t)\dot{\nu}(t') \rangle \propto \delta(t - t')$. To make progress, we

note that our measurements (the ToAs) tell us only about the average behaviour of the pulsar during a ToA gap, not the value of $d\nu/dt$ at any given instant (e.g. the end of a ToA gap). The delta-correlated nature of $\dot{\nu}(t)$ means that its value at any given instant is not correlated with its value at nearby times, and hence our measurements, which necessarily take place over a finite time span and involve an averaging process, contain no information about $\dot{\nu}(t)$ at any particular instant. The hidden state variable $\dot{\nu}$ is therefore more usefully identified as the effective rate of change in ν over a ToA gap, i.e. $z^{-1}[\nu(t_2) - \nu(t_1)]$. The dispersion in the hidden state variable $\dot{\nu}$ is then better characterized by the dispersion in ν due to the random walk, which is $\sigma_\nu z^{1/2}$, divided by z to obtain a dispersion in the average $\dot{\nu}$ over the ToA gap⁵ – hence $\text{cov}(\dot{\nu}, \dot{\nu}) = \sigma_\nu^2/z$. Under this interpretation, $\dot{\nu}$ is a linear function of $\nu(t_1)$ and $\nu(t_2)$. This presents an issue for the HMM glitch detector as currently formulated, as it implies a singular covariance matrix between ν and $\dot{\nu}$. The transition matrix $A_{q_i, q_j}(z)$ involves the inverse of this covariance matrix and is therefore undefined. To circumvent this issue, as an approximation we neglect the correlation between ν and $\dot{\nu}$. The covariances for the RWF0 model are therefore taken to be

$$\text{cov}(\nu, \nu) = \sigma_\nu^2 z, \quad (35)$$

$$\text{cov}(\dot{\nu}, \nu) = 0, \quad (36)$$

$$\text{cov}(\dot{\nu}, \dot{\nu}) = \sigma_\nu^2/z. \quad (37)$$

An alternative solution is to reformulate the HMM to track ν alone, which we defer to future work.

6.3 Online glitch detection

For every new observation which produced a good ToA, an HMM analysis of the most recent 30 ToAs⁶ for that pulsar is automatically run with low latency (~ 1 hour) to check for a potential glitch. In the event of a possible detection (any glitch candidate with a log Bayes factor greater than $10^{1/2}$; see Section 6.1), an email alert is generated and the data are further scrutinised to determine whether there is compelling evidence that a glitch has occurred, whereupon a public alert is sent out manually. In total, 2963 online HMM analyses produce a glitch candidate. In approximately, 80 per cent of cases the glitch candidate can be ascribed to RFI contamination which is not completely removed by the timing pipeline (see Section 2.3) and leads to a corrupted ToA. For 81 per cent of the remaining 539 candidates, there is no obvious contamination in the data, but further investigation – either visual inspection of the timing residuals or re-analysis using the HMM with a finer grid in ν and $\dot{\nu}$ – reveals no compelling glitch candidate, and hence no public alert is sent out. In three pulsars, clear glitches occurred while the online glitch detection pipeline was operational (approximately 2021 October through 2023 June), namely PSRs J0835–4510, J0742–2822, and J1740–3015. In all three objects the glitches were large, and easily detected by the online pipeline (Dunn et al. 2021b, 2022b, 2023b) – the remainder of the glitch candidates generated by the online pipeline contain one of these three glitches within their analysis window.

For simplicity, the online glitch searches use a uniform set of parameters for every pulsar, with values listed in Table 7. The timing

⁵This is similar to the argument in Section 5.1 to derive equation (20).

⁶The window length is chosen arbitrarily, but synthetic data tests show that glitch detection probability is a weak function of the number of ToAs; see appendix G of Melatos et al. (2020).

Table 7. HMM glitch detector parameters used in the online glitch detection analyses described in Section 6.3. Most parameters are fixed, but σ depends on the average ToA gap length in a 30-ToA window, denoted by \bar{z} .

Parameter	Value
ν_\pm	1×10^{-6} Hz
ϵ_ν	1×10^{-8} Hz
$\dot{\nu}_\pm$	2×10^{-14} Hz s ⁻¹
$\epsilon_{\dot{\nu}}$	5×10^{-16} Hz s ⁻¹
Timing noise model	RWF1
$\sigma_{\dot{\nu}}$	$\max(\epsilon_{\dot{\nu}} \bar{z}^{-1/2}, 10^{-21} \text{ Hz s}^{-3/2})$

noise model is fixed to be RWF1. The one parameter which is allowed to vary is $\sigma_{\dot{\nu}}$, which depends on the average ToA gap length within the 30-ToA window, denoted by \bar{z} .

The procedure adopted in the offline glitch searches described in Section 6.4 and Appendix C was not in hand at the time that the online pipeline was in operation. This procedure allows the offline search to achieve better sensitivity to small glitches. Estimates of the sensitivity of both the online and offline pipelines are discussed in Section 6.6.

6.4 Offline glitch detection

The offline glitch searches are conducted on the same combined EW-NS data sets analysed in Section 3. In contrast to earlier HMM analyses involving a large number of objects (Lower et al. 2021; Dunn et al. 2022a), here we make explicit use of the timing noise information returned by the ENTERPRISE analyses of Section 3 to better match the timing noise model included in the HMM with the observed data. The timing noise information returned by ENTERPRISE consists of the model selection results, and the values of A_{red} and β in equation (3), if applicable. The two models available are RWF0 and RWF1, discussed in Section 6.2.2. The models generate phase residual PSDs with spectral indices of -4 and -6 , respectively. Based on the value of β returned by ENTERPRISE, we assign a timing noise model to each pulsar, and based on the value of A_{red} we set the amplitude of the white noise driving term. The detailed procedure for choosing HMM parameters is described in Appendix C.

A number of glitches are already identified prior to running the offline glitch search, by the online search (see Section 6.3), by eye during the initial timing process using TEMPO2 and ENTERPRISE, or by other authors, e.g. Lower et al. (2020), Keith et al. (2024), and Zubieta et al. (2024). When the glitch is large ($\Delta\nu \gtrsim 10^{-7}$ Hz), the data sets passed to the HMM pipeline are split into pre- and post-glitch sections to avoid the need for a large ν range to accommodate the glitch. In one case, PSR J1048–5832, we divide the full data set in two equal halves not because of the presence of a glitch but because the strong timing noise makes an analysis of the full data set prohibitively expensive. In this case, the two halves overlap by 60 d in order to ensure that a glitch close to the boundary between the two subsets is not missed. Data sets containing smaller glitches are analysed in full, in order to check with the HMM that these by-eye identifications are not timing noise fluctuations which have been miscategorized.

The complete list of glitch candidates identified by the offline HMM pipeline is given in Table 8.

When multiple glitch candidates in a single pulsar are returned, we list only the most significant candidate. These candidates are

Table 8. Glitch candidates identified by the offline HMM glitch search described in Section 6.4. The detection threshold is $\ln K_{\text{th}} = 1.15$. Note that large ($\Delta\nu \gtrsim 10^{-7}$ Hz) glitches are not included here, as the data sets for those pulsars are broken into pre- and post-glitch sections to allow for smaller DOIs and hence more sensitive searches. In the last two columns, we indicate whether the glitch was previously identified, either as part of the UTMOST-NS online glitch detection pipeline or through by-eye inspection of the timing residuals, and whether the glitch candidate is vetoed by the procedure described in Section 6.4.

PSR	Glitch epoch (MJD)	$\ln K_1^*$	Previously identified?	Vetoed?
J0742–2822	59857–59863	1.5	N	Y
J0908–4913	60092–60119	3	N	Y
J1136+1551	57953–58098	1.3	N	Y
J1136–5525	60078–60115	2.9	N	Y
J1257–1027	58666–58697	2.7×10^2	Y	N
J1359–6038	58699–58748	4.5	N	Y
J1453–6413	59327–59359	2.1×10^2	Y	N
J1703–4851	58564–58668	5×10^2	Y	N
J1709–4429	58522–58554	2.2	N	Y
J1731–4744	58103–58140	2.3	N	Y
J1752–2806	58414–58519	1.5	N	Y
J1836–1008	59427–59458	1.9×10^3	Y	N
J1902+0615	58624–58655	16	Y	N

Table 9. Timing properties of all glitches analysed in this work. Parameter estimates are from ENTERPRISE analyses, uncertainties are 95 per cent credible intervals. The glitches are discussed in detail in Section 6.5.

PSR	Glitch epoch (MJD)	$\Delta\nu_{\text{p}}/\nu \times 10^9$	$\Delta\dot{\nu}_{\text{p}}/\dot{\nu} \times 10^3$	$\Delta\nu_{\text{d}}/\nu \times 10^9$	τ_{d} (d)
J0742–2822	59840.0	$4283.7^{+1.8}_{-2.1}$	$10.3^{+4.5}_{-4.6}$	$16.1^{+1.4}_{-1.3}$	$14.3^{+2.6}_{-2.4}$
J0835–4510	57734.4855	1432.4 ± 0.4	$6.9^{+0.4}_{-0.5}$	$9.3^{+1.7}_{-1.5}$	$1.9^{+0.7}_{-0.6}$
J0835–4510	58515.0	$2479.4^{+2.0}_{-2.2}$	7.1 ± 0.8	$17.6^{+2.4}_{-2.0}$	$7.7^{+2.3}_{-2.0}$
J0835–4510	59417.2108	1236.9 ± 0.3	5.7 ± 0.3	$6.7^{+1.3}_{-1.2}$	$2.9^{+0.8}_{-0.6}$
J0908–4913	58765.06	22.0 ± 0.2	$0.8^{+0.8}_{-0.9}$	–	–
J1105–6107	58581.0	$1171.5^{+0.4}_{-0.5}$	$2.8^{+1.9}_{-2.0}$	–	–
J1257–1027	58649.3	3.07 ± 0.03	$0.3^{+3.2}_{-3.1}$	–	–
J1453–6413	59326.0	0.9 ± 0.2	$0.0^{+0.5}_{-0.4}$	–	–
J1703–4851	58522.0	11.4 ± 0.1	5.3 ± 1.7	–	–
J1709–4429	58178.0	$2407.9^{+8.5}_{-8.2}$	$3.9^{+0.5}_{-0.7}$	$26.3^{+8.1}_{-7.7}$	$71.4^{+21.8}_{-21.6}$
J1731–4744	57984.260363	$3146.5^{+0.9}_{-0.8}$	1.5 ± 0.2	2.8 ± 0.9	$79.2^{+41.2}_{-40.5}$
J1740–3015	57468.0	229.1 ± 1.3	$1.0^{+0.9}_{-1.0}$	–	–
J1740–3015	58241.0	$837.7^{+1.2}_{-1.7}$	1.7 ± 0.7	$16.7^{+31.7}_{-16.7}$	$24.7^{+56.7}_{-24.7}$
J1740–3015	59945.0	327.9 ± 2.0	1.9 ± 1.6	–	–
J1803–2137	58958.0	$4648.2^{+37.7}_{-36.9}$	6.6 ± 1.2	–	–
J1836–1008	59400.0	$35.4^{+3.7}_{-4.1}$	$2.2^{+7.4}_{-8.4}$	–	–
J1902+0615	58590.0	0.8 ± 0.2	$0.5^{+1.9}_{-2.1}$	–	–

subjected to a simple veto procedure, whereby the data set containing the candidate is re-analysed with the two ToAs bracketing the glitch candidate removed (Dunn et al. 2022a). The veto is designed to exclude the case where a single ToA is displaced from the overall trend for non-astrophysical reasons (e.g. RFI contamination and clock errors), thereby generating a glitch candidate in the ToA gap immediately before or after the displaced ToA. If the re-analysis fails to return a candidate above the detection threshold, the initial candidate is vetoed. After vetos, the offline pipeline does not produce any additional glitch candidates over those identified by the online pipeline and by-eye inspection of the timing residuals. However it does produce tighter upper limits on the size of undetected glitches than the online pipeline, as discussed in Section 6.6.

6.5 Glitches in the sample

In addition to those glitches detected during the UTMOST-NS timing programme and those already known from previous analyses, several glitches are identified when inspecting the timing residuals by eye to prepare for the ENTERPRISE analysis of the combined data sets.

The full list of the 17 glitches known in the combined EW + NS data set is given in Table 9. The glitch in PSR J1902+0615 is the only one which has not previously been identified in the literature. The glitch parameter estimates and their 95 per cent credible intervals are derived from the ENTERPRISE analyses. To determine whether there is support for a decaying component, we perform the model selection procedure outlined in Section 3.2 twice for each glitch: once with a $\Delta\nu_{\text{d}}$ prior which is log-uniform on $[10^{-10}$ Hz, 10^{-4} Hz] and again with $\Delta\nu_{\text{d}}$ fixed at zero. If the log Bayes factor between the favoured

model with a floating $\Delta\nu_d$ and the $\Delta\nu_d = 0$ model exceeds 5, we adopt the model with a decaying component, and report constraints on the decay parameters. Previous analyses have reported evidence for multiple exponential components with different time-scales (e.g. McCulloch et al. 1987; Wong, Backer & Lyne 2001; Dodson et al. 2002; Shannon et al. 2016; Liu et al. 2024b). A comprehensive search for additional decaying components is beyond the scope of this work.

In the remainder of this section, we discuss each of the detected glitches in detail.

6.5.1 PSR J0742–2822

A new glitch was reported in PSR J0742–2822 by Shaw et al. (2022) in 2022 September, and soon after confirmed by UTMOST-NS (Dunn et al. 2022b) as well as Grover et al. (2022) and Zubieta et al. (2022a). It is confidently detected by the online pipeline in the first post-glitch observation at UTMOST-NS, taken on 22 September (MJD 59844.9), with $\ln K_1^* = 123$. With a permanent frequency step of $\Delta\nu_p/\nu = (4284 \pm 2) \times 10^{-9}$, it is by far the largest glitch out of the nine known in this pulsar; the previous largest glitch had an amplitude of $\Delta\nu/\nu = 92(2) \times 10^{-9}$ (Espinoza et al. 2011). We detect a decaying component ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 99$) with $\Delta\nu_d/\nu = 16.1_{-1.3}^{+1.4} \times 10^{-9}$ and $\tau_d = 14_{-2}^{+3}$ d. Zubieta et al. (2024) also report evidence for a decaying component, although with a somewhat longer time-scale, viz. $\Delta\nu_d/\nu = (19.16 \pm 0.04) \times 10^{-9}$ and $\tau_d = 33.4 \pm 0.5$ d.

6.5.2 PSR J0835–4510

The most recent large glitch in PSR J0835–4510 is the first one to be detected during the UTMOST-NS observing run, occurring in 2021 July. It was first reported by Sosa-Fiscella et al. (2021), and subsequently confirmed by a number of other observatories (Olney 2021; Dunn et al. 2021b; Singha et al. 2021b). Its characteristics are typical for Vela glitches, with a large permanent frequency jump of $\Delta\nu_p/\nu = (1236.9 \pm 0.3) \times 10^{-9}$. We find strong evidence for a decaying component ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 136$), with $\Delta\nu_d/\nu = (6.7_{-1.2}^{+1.3}) \times 10^{-9}$ and $\tau_d = 2.9_{-0.6}^{+0.8}$ d. Zubieta et al. (2024) reported three distinct recovery components, with time-scales of 0.994(8), 6.400(2), and 535(8) d. The reported fractional amplitudes of the two shortest time-scale recoveries are $9(1) \times 10^{-9}$ and $3(1) \times 10^{-9}$, respectively, while the long time-scale recovery has a larger reported amplitude of 512×10^{-9} . Our detected decaying component aligns roughly with the two short time-scale components, but does not fall within the error bars of either of them.

Apart from the 2021 July glitch, two earlier Vela glitches are also covered by our data set and have been discussed by a number of previous authors (Sarkissian et al. 2017; Palfreyman et al. 2018; Kerr 2019; Lopez Armengol et al. 2019; Sarkissian et al. 2019; Lower et al. 2020; Gügerciñoğlu et al. 2022). In the case of the glitch at MJD 57734, we detect a significant ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 219$) decaying component with a large amplitude of $\Delta\nu_d/\nu = 9.3_{-1.5}^{+1.7} \times 10^{-9}$ and a short time-scale of $\tau_d = 1.9_{-0.6}^{+0.7}$ d. This is approximately consistent with the analysis of Sarkissian et al. (2017), who found $\Delta\nu_d/\nu = (11.5 \pm 0.7) \times 10^{-9}$ and $\tau_d = (0.96 \pm 0.17)$ d. We also find strong evidence for a recovery in the glitch at MJD 58515 ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 86.3$), with $\Delta\nu_d/\nu = 18 \pm 2 \times 10^{-9}$ and $\tau_d = 7.7_{-2.0}^{+2.3}$ d.

Previous Vela glitches have been observed to recover on a wide range of time-scales, from minutes to hundreds of days (McCulloch et al. 1990; Shannon et al. 2016; Ashton et al. 2019a), with several previous studies reporting an intermediate recovery time-scale on the order of days (Flanagan 1990; McCulloch et al. 1990; Dodson et al.

2002; Sarkissian et al. 2017). In all three glitches analysed here we find strong evidence for a recovering component on this intermediate time-scale and with amplitudes $\sim 10^{-7}$ Hz, similar to previous works. The recovery fractions $Q = \Delta\nu_d/(\Delta\nu_p + \Delta\nu_d)$ for the three glitches in our data set are (from oldest to newest) 6.5×10^{-3} , 7.1×10^{-3} , and 5.3×10^{-3} , compared to 6.5×10^{-3} as reported by McCulloch et al. (1990) and 5.5×10^{-3} as reported by Dodson et al. (2002). An intermediate-time-scale decaying component with a recovery fraction of approximately 5×10^{-3} appears to be a common feature of large Vela glitches.

6.5.3 PSR J0908–4913

This glitch was previously detected and discussed by Lower et al. (2019). It is the only glitch known in PSR J0908–4913. Our recovered amplitude of $\Delta\nu_p/\nu = (22.0 \pm 0.2) \times 10^{-9}$ is only marginally different from the value inferred by Lower et al. (2019) of $\Delta\nu_p/\nu = (21.7 \pm 0.1) \times 10^{-9}$. Like Lower et al. (2019), we are unable to distinguish $\Delta\dot{\nu}_p$ from zero within the measurement uncertainty. Lower et al. (2019) did not report an exponential recovery, and we find no evidence for a decaying component.

6.5.4 PSR J1105–6107

This glitch occurred in 2019 April, between MJD 58580 and MJD 58584. It occurred while UTMOST-EW observations were still ongoing – the last UTMOST-EW observation of this pulsar was taken on 2020 April 28. It was not reported in previous analyses (Lower et al. 2020; Dunn et al. 2022a) because it came after the cut-off for the first data release. However, a measurement was reported by Abbott et al. (2022b) based on UTMOST-EW data, as part of a search for continuous gravitational wave emission in the aftermath of the glitch. Although there are no UTMOST-NS ToAs available for this pulsar, here we re-analyse the UTMOST-EW data covering this glitch to provide updated parameter and uncertainty estimates. We find $\Delta\nu_p/\nu = 1171.4_{-0.5}^{+0.4} \times 10^{-9}$ and $\Delta\dot{\nu}_p/\dot{\nu} = 2.8_{-2.0}^{+1.9} \times 10^{-3}$. This is the largest glitch reported in PSR J1105–6107; the next largest has a magnitude of $\Delta\nu_p/\nu = (971.7 \pm 0.5) \times 10^{-9}$ (Yu et al. 2013). As in Abbott et al. (2022b), we do not find evidence for a decaying component.

6.5.5 PSR J1257–1027

This glitch was reported by Lower et al. (2020) and is the only known glitch in this pulsar. Our recovered amplitude of $\Delta\nu_p/\nu = (3.07 \pm 0.03) \times 10^{-9}$ is consistent with the Lower et al. (2020) value of $\Delta\nu_p/\nu = 3.20_{-0.57}^{+0.16} \times 10^{-9}$. Lower et al. (2020) reported an upper limit on the change in frequency derivative, $\Delta\dot{\nu}_p/\dot{\nu} \lesssim 286 \times 10^{-3}$. With our extended data set, we measure $\Delta\dot{\nu}_p/\dot{\nu} = 0.3_{-3.1}^{+3.2} \times 10^{-3}$. We are unable to distinguish $\Delta\dot{\nu}_p$ from zero. Lower et al. (2020) did not detect an exponential recovery, and we find no evidence for a decaying component.

6.5.6 PSR J1453–6413

This glitch is identified by eye as occurring during the gap between the UTMOST-EW and UTMOST-NS data sets, between MJDs 58927 and 59326. It was subsequently confirmed in the offline glitch search, being confidently detected with $\ln K_1^* = 2.1 \times 10^2$. It was reported by Li et al. (2023) as occurring at MJD 59060(12) with $\Delta\nu/\nu = 1.180(7) \times 10^{-9}$, and by Keith et al. (2024) as occurring at MJD

59015 with $\Delta\nu/\nu = (1.14 \pm 0.13) \times 10^{-9}$. Our value of $\Delta\nu_p/\nu = (0.9 \pm 0.2) \times 10^{-9}$ is in mild tension with these values, and our limit on $\Delta\dot{\nu}_p/\dot{\nu} \lesssim 0.5 \times 10^{-3}$ is marginally consistent with the value measured by Li et al. (2023) of $0.50(3) \times 10^{-3}$; Keith et al. (2024) did not report a measurement of $\Delta\dot{\nu}$. It is likely that a combination of the long observing gap in the UTMOST data and differences in timing noise modelling account for these tensions. Li et al. (2023) employ a method based on a Cholesky decomposition (Coles et al. 2011; Dang et al. 2020) to model the red noise in the data set, rather than the Fourier decomposition used in this work. We remind the reader that the results of any glitch search are conditional on the phase model and indeed the user-selected false alarm and false dismissal probabilities.

6.5.7 PSR J1703–4851

This glitch was first reported by Lower et al. (2020) and is the only known glitch in this pulsar. Our recovered glitch amplitude is somewhat smaller than the value reported by Lower et al. (2020), viz. $\Delta\nu_p/\nu = (11.4 \pm 0.1) \times 10^{-9}$ rather than $\Delta\nu_p/\nu = 19.0_{-0.7}^{+1.0} \times 10^{-9}$. We also find a significantly different value of $\Delta\dot{\nu}/\dot{\nu} = (5.3 \pm 1.7) \times 10^{-3}$ as compared to the Lower et al. (2020) value of $292_{-53}^{+38} \times 10^{-3}$. The updated measurement of $\Delta\dot{\nu}$ is largely enabled by the addition of the NS data – the EW data used in Lower et al. (2020) are quite sparse post-glitch, with only four post-glitch ToAs available, covering six months. Future timing is likely to further refine this measurement, as the timing noise in this pulsar is low, so measurements of the glitch parameters are not confounded by timing noise as the timing baseline is lengthened.

6.5.8 PSR J1709–4429

The glitch at MJD 58 178 was first reported by Lower et al. (2018), with $\Delta\nu/\nu = (52.4 \pm 0.1) \times 10^{-9}$. This glitch amplitude was later discovered to be in error in the course of a re-analysis using the HMM glitch detector, and revised to $\Delta\nu/\nu = (2432.2 \pm 0.1) \times 10^{-9}$ by Dunn, Lower & Melatos (2021a). Although Dunn et al. (2021a) found no evidence for a decaying component, here we find strong evidence ($\ln \mathcal{B}_{\nu_d \neq 0}^{\Delta\nu_d} = 42$) for a decaying term with $\Delta\nu_d/\nu = 26.3_{-7.7}^{+8.1} \times 10^{-9}$ and $\tau_d = 71 \pm 22$ d. The discrepancy may be due to the difference in the available post-glitch data – Dunn et al. (2021a) used data from the first UTMOST data release (Lower et al. 2020) which extends up to MJD 58418, 240 d after the glitch, while this analysis includes 710 d of post-glitch data before a long observing gap in early 2020. The total glitch amplitude $(\Delta\nu_p + \Delta\nu_d)/\nu = (2434 \pm 10) \times 10^{-9}$ is consistent with the value reported by Dunn et al. (2021a). The fractional change in frequency derivative reported by Dunn et al. (2021a) of $\Delta\dot{\nu}_p/\dot{\nu} = (4.7 \pm 0.3) \times 10^{-3}$ is marginally consistent with the value obtained in this analysis, viz. $3.9_{-0.7}^{+0.5} \times 10^{-3}$.

6.5.9 PSR J1731–4744

The glitch at MJD 57984 was first reported by Jankowski et al. (2017), and is the largest glitch ever observed in this pulsar. None of the previous works analysing this glitch (Jankowski et al. 2017; Basu et al. 2020; Lower et al. 2020) have detected a post-glitch recovery. In our analysis, we find moderate evidence ($\ln \mathcal{B}_{\Delta\nu_d \neq 0}^{\Delta\nu_d} = 6.4$) for the presence of a decaying component with $\Delta\nu_d/\nu = 2.8 \pm 0.9 \times 10^{-9}$ and $\tau_d = 79 \pm 41$ d. It is possible that our extended data set relative to other studies allows for better characterization of the recovery. However, given that there is significant timing noise in this pulsar

($\log_{10} A_{\text{red}} = -9.7, \beta = 3.7$) it is also possible that there is confusion between the timing noise and a putative long time-scale recovery process.

6.5.10 PSR J1740–3015

PSR J1740–3015 most recently suffered a glitch in 2022 December, first reported by Zubieta et al. (2022b) and subsequently confirmed by Grover et al. (2023) and UTMOST-NS (Dunn et al. 2023b). It was confidently detected by the online pipeline in the first post-glitch observation at UTMOST-NS, taken on 2023 January 1 (MJD 59945), with $\ln K_1^* = 79$. It is similar to other glitches observed in this pulsar, with a moderate amplitude of $\Delta\nu_p/\nu = (327 \pm 2) \times 10^{-9}$ and a change in frequency derivative of $\Delta\dot{\nu}/\dot{\nu} = 1.9 \pm 1.6 \times 10^{-3}$. No recovery is observed for this glitch. Zubieta et al. (2024) reported a recovering component with $\Delta\nu_d/\nu = 14.59(3) \times 10^{-9}$ and $\tau_d = 124(2)$ d. It is possible that we do not detect a recovering component is because we have only 179 d of post-glitch timing data for this event, i.e. less than two e-foldings if the decay time-scale reported by Zubieta et al. (2024) is accurate.

There are several previously reported glitches within our data set for this pulsar. Two, reported as occurring at MJDs 57296.5 (Lower et al. 2020) and 57346 (Jankowski et al. 2015; Basu et al. 2020; Lower et al. 2020), are small, with $\Delta\nu/\nu \lesssim 2 \times 10^{-9}$. We do not include these glitches in our timing model here. They are not detected by the HMM offline glitch search (note they lie below the 90 per cent upper limit $\Delta\nu^{90 \text{ per cent}}/\nu = 1.4 \times 10^{-8}$ for the relevant stretch of data; see Section 6.6) and visual inspection of the residuals does not reveal any obvious glitch signature. The other two glitches, at MJDs 57468 and 58241, are similar in size to the latest glitch and are clearly visible. We recover glitch parameters which are consistent with previous studies (Jankowski et al. 2016; Basu et al. 2020; Lower et al. 2020; Liu et al. 2024b). We find moderate evidence for a recovery term in the glitch at MJD 58241 ($\ln \mathcal{B}_{\Delta\nu_d \neq 0}^{\Delta\nu_d} = 6.4$), with $\Delta\nu_d/\nu = 16.7_{-16.7}^{+31.7}$ and $\tau_d = 25_{-25}^{+57}$. This recovery was not reported by other analyses (Basu et al. 2020; Lower et al. 2020; Liu et al. 2024b).

6.5.11 PSR J1803–2137

This glitch is identified by eye during the initial timing process as occurring during the gap between the EW and NS data sets, and is a large glitch, with $\Delta\nu/\nu = 4648_{-37}^{+38} \times 10^{-9}$ and $\Delta\dot{\nu}/\dot{\nu} = 6.6 \pm 1.2 \times 10^{-3}$. This glitch was previously reported by Keith et al. (2024) as occurring at MJD 58920 with $\Delta\nu_p/\nu = (4702 \pm 11) \times 10^{-9}$. This value is in moderate tension with our measurement, which we ascribe to the glitch occurring in a 700-d observing gap. This glitch has also been recorded in the Jodrell Bank Observatory (JBO) glitch catalogue⁷ (Basu et al. 2022) as occurring at MJD 58958 \pm 40, with $\Delta\nu/\nu = (4960.6 \pm 1.3) \times 10^{-9}$ and $\Delta\dot{\nu}/\dot{\nu} = (8.50 \pm 0.3) \times 10^{-3}$ – different again from the values reported here and by Keith et al. (2024).

6.5.12 PSR J1836–1008

This glitch is identified by eye as occurring during the gap between the EW and NS data sets, and was subsequently confirmed in the offline glitch search with $\ln K_1^* = 2.1 \times 10^2$ with an estimated glitch epoch between MJD 59427 and MJD 59458. It is the second glitch

⁷<https://www.jb.man.ac.uk/pulsar/glitches/gTable.html>

reported in this pulsar. With an estimated $\Delta\nu/\nu = 35 \pm 4 \times 10^{-9}$, it is an order of magnitude larger than the previous known glitch, which had $\Delta\nu/\nu = 3.6(1) \times 10^{-9}$ (Basu et al. 2022). No significant change in $\dot{\Delta\nu}/\dot{\nu}$ is detected. It was also detected by Keith et al. (2024) as occurring at MJD 58950, with $\Delta\nu/\nu = (32.8 \pm 0.9) \times 10^{-9}$. Like previous authors, we find no evidence for a decaying component.

6.5.13 PSR J1902+0615

This glitch is identified by eye as occurring between MJD 58589 and MJD 58595, in 2019 April. It was subsequently confirmed in the offline glitch search, with $\ln K_1^* = 7.2$. It has not been reported elsewhere. Unlike the other glitches detected in the combined data set, this glitch did not occur in the observing gap – the UTMOST-EW data for this pulsar extends up to 2020 July. However, it was not detected in previous searches (Lower et al. 2020; Dunn et al. 2022a) because the cut-off date for this pulsar in the data release analysed in those studies is 2018 November. The glitch is small, with $\Delta\nu/\nu = (0.8 \pm 0.2) \times 10^{-9}$, but it is the largest of the seven known glitches in this pulsar. Other reported glitches have fractional sizes between 2×10^{-10} and 4×10^{-10} (Yuan et al. 2010; Espinoza et al. 2011; Basu et al. 2022).

6.6 Glitch upper limits

The automated nature of the HMM glitch detector allows us to compute systematic upper limits on the size of glitches that could have been detected in a given analysis. This involves first creating fake data sets which mimic the real data sets, i.e. they have the same spin parameters, and white and red noise injected using LIBSTEMPO at the levels recovered by ENTERPRISE. Glitches of various sizes are injected into these fake data sets, and the detection probability as a function of glitch size, denoted $P_d(\Delta\nu)$, is estimated by running the HMM on the fake data sets and checking in each instance whether a glitch candidate is recovered. We take the benchmark detection probability to be 90 per cent, and denote the glitch size at which 90 per cent of injected glitches are successfully detected by $\Delta\nu^{90 \text{ per cent}}$. Previous sensitivity estimates of the HMM have employed a binary search method to determine $\Delta\nu^{90 \text{ per cent}}$ (Lower et al. 2021; Dunn et al. 2022a, 2023a). In this work, we adopt a technique borrowed from sensitivity estimation in the context of continuous gravitational wave searches (Banagiri et al. 2019; Abbott et al. 2021, 2022a) and assume the detection probability is a sigmoid in $\log_{10} \Delta\nu$ of the form

$$P_d(\Delta\nu) = [1 + \exp\{-r[\log_{10}(\Delta\nu/\text{Hz}) - c_0]\}]^{-1}. \quad (38)$$

We fit for the dimensionless parameters r and c_0 using the LMFIT package (Newville et al. 2014), and invert the best-fitting sigmoid curve to find the value of $\Delta\nu^{90 \text{ per cent}}$ satisfying $P_d(\Delta\nu^{90 \text{ per cent}}) = 0.9$.

The above procedure applies to both the online and offline glitch searches. In the online glitch searches, we randomly choose 10 30-ToA windows within the total UTMOST-NS data set and compute the value of $\Delta\nu^{90 \text{ per cent}}$ for each. The injected glitch epochs are chosen uniformly within the window. We expect some variation in these values as the observing cadence and telescope sensitivity (and hence the ToA uncertainties) both vary by approximately a factor of two over the two years of operation (see Appendix A). For each pulsar, we take the average over the set of chosen windows and quote a mean $\Delta\nu^{90 \text{ per cent}}$ value. The distribution of $\Delta\nu^{90 \text{ per cent}}/\nu$ is shown in the top panel Fig. 8, with a mean $\Delta\nu^{90 \text{ per cent}}/\nu$ of 4.7×10^{-8} .




Figure 8. Distribution of 90 per cent confidence upper limits on fractional glitch size for the online (top) and offline (bottom) glitch searches described in Sections 6.3 and 6.4 ($\Delta\nu^{90\%}/\nu$), compared to the observed distribution of fractional glitch sizes recorded in the JBO glitch catalogue. All histograms are normalized so that they integrate to one. The online upper limits include one value per pulsar, 158 total. The offline upper limits include one value per time span analysed per pulsar. When a data set has been split into pre- and post-glitch sections or has been subdivided because of strong timing noise, there are multiple time spans per pulsar – in total, there are 173 offline upper limits. The JBO values include all glitches with positive $\Delta\nu$, 676 in total. In the case of the online search, the mean upper limit across the sample is $\Delta\nu^{90 \text{ per cent}}/\nu = 4.6 \times 10^{-8}$ with a maximum of 5.0×10^{-7} (PSR J1703–4851) and a minimum of 6.8×10^{-11} (PSR J1939+2134). In the case of the offline search, the mean upper limit is 6.3×10^{-9} , with a maximum of 1.3×10^{-7} (PSR 1703–1846) and a minimum of 7.2×10^{-12} (PSR J0437–4715).

This is comparable to the sensitivity achieved in the offline search of the UTMOST-EW data set conducted by Dunn et al. (2022a), where the mean $\Delta\nu^{90 \text{ per cent}}/\nu$ value is 1.9×10^{-8} .

The upper limits in the offline search case are computed in the same way, except that the complete EW+NS data sets are used and the injected glitch epochs are chosen uniformly within the total time span of each data set. In cases where the EW + NS data set for a pulsar is divided by the presence of large glitches, we derive separate upper limits for each section. The bottom panel of Fig. 8 shows the distribution of $\Delta\nu^{90 \text{ per cent}}/\nu$ for the offline searches. Compared to the equivalent distribution in Dunn et al. (2022a), as well as the upper limits for the online pipeline as shown in the top panel of the same figure, the distribution is shifted towards smaller values,

with a mean of 6.3×10^{-9} compared to 1.9×10^{-8} in Dunn et al. (2022a) and 4.7×10^{-8} for the online glitch searches. The minimum upper limit for the offline search is 7.2×10^{-12} , achieved for PSR J0437–4715, compared to the minimum 6.8×10^{-11} achieved for PSR J1939+2134 in the online case. Similarly the maximum upper limit in the offline search is 1.4×10^{-7} , achieved for PSR J1703–1846, compared to the maximum 5.0×10^{-7} achieved for PSR J1703–4851 in the online case.

In the online searches, the primary factors determining sensitivity are the typical ToA uncertainty and observing cadence (Melatos et al. 2020; Dunn et al. 2023a). As discussed in Section 6.3, for the online searches we adopt a fixed configuration of the HMM glitch detector for all pulsars, and do not incorporate information about the characteristics of timing noise in the pulsar. In contrast, in the offline searches, information about the timing noise is used (see Appendix C), so that the intrinsic properties of the pulsar also play a role in determining the sensitivity; see section 6 of Melatos et al. (2020).

Small glitches have recently been reported in PSR J1048–5832 (Liu et al. 2024a; Zubieta et al. 2023, 2024) and PSR J0835–4510 (Zubieta et al. 2024) which are within the time span covered by the EW + NS data set. The four reported glitches in PSR J1048–5832 have fractional amplitudes ranging between 0.91×10^{-9} and 9.9×10^{-9} , while the small glitch in PSR J0835–4510 reported by Zubieta et al. (2024) has a fractional amplitude of 2×10^{-10} . None of these glitches are detected by the offline HMM search. They are smaller than our 90 per cent upper limits, which are 2.9×10^{-8} for PSR J1048–5832 and 2.3×10^{-8} for PSR J0835–4510.

7 CONCLUSION

In total, the UTMOST-NS timing programme produced 269 pulsar-years of new timing baseline across 167 pulsars, and 3665 h of integrated observing time. In this paper, we describe the programme, and report some of the first scientific results drawing on the full sample of observed pulsars, with a focus on irregularities in the timing behaviour. We augment the new UTMOST-NS data with existing data from the earlier UTMOST programme using the east-west arm (Jankowski et al. 2019; Lower et al. 2020), and execute two Bayesian analysis pipelines: ENTERPRISE (Ellis et al. 2020) selects between models with different combinations of timing noise and non-zero second frequency derivative, and an HMM-based glitch detector (Melatos et al. 2020) searches for glitches and produces systematic upper limits on the size of undetected glitches for each object.

We investigate the timing noise properties across the population, and find results consistent with earlier studies (Shannon & Cordes 2010; Parthasarathy et al. 2019; Lower et al. 2020) which suggest that the pulsar characteristic age $\tau_c \propto \nu \dot{\nu}^{-1}$ is correlated with the contribution of timing noise to the total timing residuals, σ_{RN} . Assuming $\sigma_{\text{RN}} \propto \nu^a \dot{\nu}^b$, we find $a = -0.85^{+0.38}_{-0.35}$ and $b = 0.56 \pm 0.16$. We find that fixing the scaling of σ_{RN} with observing time span according to the measured noise PSD spectral index improves our ability to estimate the contribution of timing noise to the timing residuals. The scatter in the relationship between σ_{RN} and its predicted value decreases by approximately 10 per cent when the scaling with observation time span takes its theoretically expected form, and we find $a = -0.49^{+0.31}_{-0.32}$ and $b = 0.47 \pm 0.13$. We additionally test the consistency of the estimated noise PSD parameters under variations in data set and methodology, and find that the estimated parameters of the noise PSD remain consistent as additional data are incorporated, with no significant systematic shift in either the amplitude A_{red} or the

spectral index β of the PSD. Moreover, no significant shift in either A_{red} or β is observed when comparing the results of ENTERPRISE against TEMPONEST, using the same data set.

We investigate the presence of secular $\ddot{\nu}$ terms for the pulsars in our sample, finding 39 pulsars for which the $\ddot{\nu}$ posterior excludes zero at 95 per cent confidence. However, when simultaneously modelling the timing noise and $\ddot{\nu}$ term, measurements of $\ddot{\nu} \neq 0$ are still contaminated by timing noise, as previously noted by other authors (Keith & Nițu 2023; Vargas & Melatos 2023). In particular, the ensemble variation in the measured $\ddot{\nu}$ can exceed the uncertainty reported by ENTERPRISE by a typical factor of approximately three when the spectral index of the timing noise PSD is approximately six. This effect is closely connected to the low-frequency cut-off in the timing noise model. Fixing the cut-off at $1/2T_{\text{span}}$, instead of allowing for a cut-off at either $1/T_{\text{span}}$ or $1/2T_{\text{span}}$, we still recover 17 non-zero $\ddot{\nu}$ values, with the implied braking index ranging between $-1.9^{+1.1}_{-0.9} \times 10^4$ and $3.2^{+3.0}_{-2.8} \times 10^3$. If we additionally impose the restriction that a model with $\ddot{\nu}$ included must be favoured over a model with $\ddot{\nu}$ fixed at 0 by a log Bayes factor of at least 5, the only pulsar with a significant measurement of $\ddot{\nu}$ is the Crab pulsar, for which we measure $n = 2.54 \pm 0.03$. We expand on the work of Vargas & Melatos (2023) to give explicit analytic conditions under which ‘anomalous’ braking indices are expected (i.e. $|n| \gg 1$) in terms of the parameters estimated by ENTERPRISE. These conditions involve the amplitude and spectral index of the timing noise PSD, the spin frequency and its first derivative, and the observing time span. We show that for all but two pulsars in our sample, variation in the measured n due to timing noise is expected theoretically to overwhelm the underlying secular braking index n_{pl} , under the assumption that $1 \lesssim n_{\text{pl}} \lesssim 7$. The median timing baseline required to overcome this effect is 280 yr for pulsars with spectral indices close to 4, and 3.9×10^7 yr for those with spectral indices close to 6.

Finally, we report on the presence of glitches in our sample, using the HMM pipeline to perform both an ‘online’ low-latency search for glitches, which was in place during active operation, and an ‘offline’ search which was performed afterwards. We analyse 17 glitches in total, of which three were detected while the UTMOST-NS timing programme was active and one, in PSR J1902+0615, is previously unpublished. Glitch parameters are estimated using ENTERPRISE, and we check for the presence of an exponentially decaying term via Bayesian model selection. We find evidence for decaying components in seven glitches among four pulsars, with time-scales ranging between $1.9^{+0.7}_{-0.6}$ and $79.2^{+41.2}_{-40.5}$ d and recovery fractions on the order of 10^{-3} , within the typically observed ranges for large glitches (e.g. Yu et al. 2013; Lower et al. 2021; Liu et al. 2024b). For all three large glitches in PSR J0835–4510, we measure a short decay time-scale between 1 and 10 d, and a recovery fraction of approximately 6×10^{-3} , matching several previous studies (Flanagan 1990; McCulloch et al. 1990; Dodson et al. 2002; Sarkissian et al. 2017; Zubieta et al. 2024). Measurement of these decay time-scales, particularly those on the order of days, is enabled by the high cadence of the observing programme.

We also report estimates of the sensitivity of both the online and offline glitch detection pipelines, leveraging the fast and automated nature of the HMM pipeline to derive upper limits on the size of undetected glitches through the use of automated injection studies. We find that the online pipeline achieves an average sensitivity of $\Delta \nu^{90 \text{ per cent}} / \nu = 4.6 \times 10^{-8}$, while the offline pipeline achieves a significantly better average sensitivity of $\Delta \nu^{90 \text{ per cent}} / \nu = 6.3 \times 10^{-9}$ due to the explicit incorporation of information about the timing noise.

This study of bright southern pulsars is the last major project to run at the Molonglo Telescope. It is being decommissioned in 2024 after almost 60 yr of operation as a significant Australian radio astronomy facility.

ACKNOWLEDGEMENTS

Parts of this research are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project nos CE170100004 and CE230100016). LD is supported by an Australian Government Research Training Program Scholarship. This work was performed in part on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy allocation provided by the Australian Government. The Molonglo Observatory is owned and operated by the University of Sydney with support from the School of Physics and the University. The UTMOST project is also supported by the Swinburne University of Technology.

DATA AVAILABILITY

The data products underlying this work will be made available through the Pulsar Data Portal (<https://pulsars.org.au>) and are also available upon reasonable request to the corresponding author.

REFERENCES

- Abbott R. et al., 2021, *ApJ*, 921, 80
 Abbott R. et al., 2022a, *Phys. Rev. D*, 105, 022002
 Abbott R. et al., 2022b, *ApJ*, 932, 133
 Abolmasov P., Biryukov A., Popov S. B., 2024, *Galaxies*, 12, 7
 Alpar M. A., Pines D., Anderson P. W., Shaham J., 1984, *ApJ*, 276, 325
 Alpar M. A., Nandkumar R., Pines D., 1986, *ApJ*, 311, 197
 Antonelli M., Basu A., Haskell B., 2023, *MNRAS*, 520, 2813
 Ashton G., Lasky P. D., Graber V., Palfreyman J., 2019a, *Nat. Astron.*, 3, 1143
 Ashton G. et al., 2019b, *ApJS*, 241, 27
 Bailes M. et al., 2017, *Publ. Astron. Soc. Aust.*, 34, e045
 Banagiri S., Sun L., Coughlin M. W., Melatos A., 2019, *Phys. Rev. D*, 100, 024034
 Basu A., Joshi B. C., Krishnakumar M. A., Bhattacharya D., Nandi R., Bandhopadhyay D., Char P., Manoharan P. K., 2020, *MNRAS*, 491, 3182
 Basu A. et al., 2022, *MNRAS*, 510, 4049
 Blandford R. D., Romani R. W., 1988, *MNRAS*, 234, 57P
 Boynton P. E., Groth E. J., Hutchinson D. P., Nanos G. P. J., Partridge R. B., Wilkinson D. T., 1972, *ApJ*, 175, 217
 Caballero R. N. et al., 2016, *MNRAS*, 457, 4421
 Camilo F., Ransom S. M., Halpern J. P., Reynolds J., Helfand D. J., Zimmerman N., Sarkissian J., 2006, *Nature*, 442, 892
 Carlin J. B., Melatos A., 2019, *MNRAS*, 488, 4890
 Carlin J. B., Melatos A., 2020, *MNRAS*, 494, 3383
 Carlin J. B., Melatos A., Vukcevic D., 2019, *MNRAS*, 482, 3736
 Chamel N., 2017, *JA&A*, 38, 43
 Chukwude A. E., Chidi Odo F., 2016, *Res. Astron. Astrophys.*, 16, 150
 Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, *MNRAS*, 418, 561
 Cordes J. M., Downs G. S., 1985, *ApJS*, 59, 343
 Cordes J. M., Helfand D. J., 1980, *ApJ*, 239, 640
 D’Alessandro F., McCulloch P. M., Hamilton P. A., Deshpande A. A., 1995, *MNRAS*, 277, 1033
 Dang S. J. et al., 2020, *ApJ*, 896, 140
 Day C. K., 2022, PhD thesis, Swinburne University of Technology
 Deller A., Flynn C., 2020, *Nat. Astron.*, 4, 292
 Dewey R. J., Cordes J. M., 1989, in Ögelman H., van den Heuvel E. P. J. eds, NATO Advanced Study Institute (ASI) Series C Vol. 262, Timing Neutron Stars. Kluwer Academic / Plenum Publishers, New York, NY USA, p. 119
 Dodson R. G., McCulloch P. M., Lewis D. R., 2002, *ApJ*, 564, L85
 Dunn L., Lower M. E., Melatos A., 2021a, *MNRAS*, 504, 3399
 Dunn L. et al., 2021b, *Astron. Telegram*, 14807, 1
 Dunn L. et al., 2022a, *MNRAS*, 512, 1469
 Dunn L. et al., 2022b, *Astron. Telegram*, 15631, 1
 Dunn L., Melatos A., Espinoza C. M., Antonopoulou D., Dodson R., 2023a, *MNRAS*, 522, 5469
 Dunn L. et al., 2023b, *Astron. Telegram*, 15839, 1
 Edwards R. T., Hobbs G. B., Manchester R. N., 2006, *MNRAS*, 372, 1549
 Ellis J. A., Vallisneri M., Taylor S. R., Baker P. T., 2020, *ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite*, Zenodo, doi:10.5281/zenodo.4059815
 Espinoza C. M., Lyne A. G., Stappers B. W., Kramer M., 2011, *MNRAS*, 414, 1679
 Espinoza C. M., Antonopoulou D., Stappers B. W., Watts A., Lyne A. G., 2014, *MNRAS*, 440, 2755
 Eya I. O., Urama J. O., Chukwude A. E., 2019, *Res. Astron. Astrophys.*, 19, 089
 Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
 Flanagan C. S., 1990, *Nature*, 345, 416
 Fuentes J. R., Espinoza C. M., Reisenegger A., Shaw B., Stappers B. W., Lyne A. G., 2017, *A&A*, 608, A131
 Fuentes J. R., Espinoza C. M., Reisenegger A., 2019, *A&A*, 630, A115
 Garthwaite M. C., Hill A. R., Yates S., Twilley B., 2013, Geodetic Survey of the Molonglo Radio Observatory, Technical Report, Canberra Geoscience Australia, Canberra, Australia, <https://ecat.ga.gov.au/geonetwork/srv/en/catalog.search#/metadata/75653>
 Goncharov B., Zhu X.-J., Thrane E., 2020, *MNRAS*, 497, 3264
 Goncharov B. et al., 2021, *MNRAS*, 502, 478
 Grover H., Singha J., Joshi B. C., Arumugam P., 2022, *Astron. Telegram*, 15629, 1
 Grover H. et al., 2023, *Astron. Telegram*, 15851, 1
 Gügürcinoğlu E., Ge M. Y., Yuan J. P., Zhou S. Q., 2022, *MNRAS*, 511, 425
 Haskell B., Melatos A., 2015, *Int. J. Mod. Phys. D*, 24, 1530008
 Haskell B., Sedrakian A., 2018, in Rezzolla L., Pizzochero P., Jones D. I., Rea N., Vidiña I. eds, The Physics and Astrophysics of Neutron Stars. Astrophysics and Space Science Library Vol. 457, Springer Nature Switzerland AG, Cham, Switzerland, p. 401
 Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, *A&AS*, 47, 1
 Ho W. C. G. et al., 2020, *MNRAS*, 498, 4605
 Hobbs G., Lyne A. G., Kramer M., Martin C. E., Jordan C., 2004, *MNRAS*, 353, 1311
 Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655
 Hobbs G., Lyne A. G., Kramer M., 2010, *MNRAS*, 402, 1027
 Hotan A. W., van Straten W., Manchester R. N., 2004, *Publ. Astron. Soc. Aust.*, 21, 302
 Jankowski F. et al., 2015, *Astron. Telegram*, 8298, 1
 Jankowski F. et al., 2016, *Astron. Telegram*, 9054, 1
 Jankowski F. et al., 2017, *Astron. Telegram*, 10770, 1
 Jankowski F., van Straten W., Keane E. F., Bailes M., Barr E. D., Johnston S., Kerr M., 2018, *MNRAS*, 473, 4436
 Jankowski F. et al., 2019, *MNRAS*, 484, 3691
 Jennings R. J., Nanograv Collaboration, Chime/Pulsar Collaboration, 2022, *Astron. Telegram*, 15223, 1
 Johnston S., Galloway D., 1999, *MNRAS*, 306, L50
 Johnston S. et al., 2020, *MNRAS*, 493, 3608
 Jones P. B., 1990, *MNRAS*, 246, 364
 Jones P. B., 1998, *MNRAS*, 296, 217
 Keane E. F., 2010, PhD thesis, University of Manchester, UK
 Keith M. J., Nițu I. C., 2023, *MNRAS*, 523, 4603
 Keith M. J. et al., 2024, *MNRAS*, 530, 1581
 Kerin A. D., Melatos A., 2022, *MNRAS*, 514, 1628
 Kerr M., 2019, *Astron. Telegram*, 12481, 1
 Kumamoto H. et al., 2021, *MNRAS*, 501, 4490

- Lentati L., Alexander P., Hobson M. P., Taylor S., Gair J., Balan S. T., van Haasteren R., 2013, *Phys. Rev. D*, 87, 104021
- Levin L. et al., 2010, *ApJ*, 721, L33
- Levin L. et al., 2019, *MNRAS*, 488, 5251
- Li W. et al., 2023, *Res. Astron. Astrophys.*, 23, 105014
- Link B. K., Epstein R. I., 1991, *ApJ*, 373, 592
- Link B., Epstein R. I., 1996, *ApJ*, 457, 844
- Liu K., Verbiest J. P. W., Kramer M., Stappers B. W., van Straten W., Cordes J. M., 2011, *MNRAS*, 417, 2916
- Liu P. et al., 2024a, *MNRAS*, 533, 4274
- Liu Y. et al., 2024b, *MNRAS*, 532, 859
- Lopez Armengol F. G. et al., 2019, *Astron. Telegram*, 12482, 1
- Lower M. E. et al., 2018, *Res. Notes Am. Astron. Soc.*, 2, 139
- Lower M. E. et al., 2019, *Res. Notes Am. Astron. Soc.*, 3, 192
- Lower M. E. et al., 2020, *MNRAS*, 494, 228
- Lower M. E. et al., 2021, *MNRAS*, 508, 3251
- Lyne A., Graham-Smith F., 2012, *Pulsar Astronomy*, Cambridge University Press, Cambridge, UK
- Lyne A. G., Shemar S. L., Smith F. G., 2000, *MNRAS*, 315, 534
- Lyne A. G., Jordan C. A., Graham-Smith F., Espinoza C. M., Stappers B. W., Weltevrede P., 2015, *MNRAS*, 446, 857
- Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, *AJ*, 129, 1993
- Mandlik A., 2024, PhD thesis, Swinburne University of Technology
- Mandlik A. et al., 2024, *MNRAS*, 532, 2644
- Maron O., Serylak M., Kijak J., Krzeszowski K., Mitra D., Jessner A., 2013, *A&A*, 555, A28
- Matsakis D. N., Taylor J. H., Eubanks T. M., 1997, *A&A*, 326, 924
- McCulloch P. M., Klekociuk A. R., Hamilton P. A., Royle G. W. R., 1987, *Aust. J. Phys.*, 40, 725
- McCulloch P. M., Hamilton P. A., McConnell D., King E. A., 1990, *Nature*, 346, 822
- Melatos A., Link B., 2014, *MNRAS*, 437, 21
- Melatos A., Peralta C., Wyithe J. S. B., 2008, *ApJ*, 672, 1103
- Melatos A., Howitt G., Fulgenzi W., 2018, *ApJ*, 863, 196
- Melatos A., Dunn L. M., Suvorova S., Moran W., Evans R. J., 2020, *ApJ*, 896, 78
- Meyers P. M., Melatos A., O'Neill N. J., 2021, *MNRAS*, 502, 3113
- Middleditch J., Marshall F. E., Wang Q. D., Gotthelf E. V., Zhang W., 2006, *ApJ*, 652, 1531
- Millhouse M., Melatos A., Howitt G., Carlin J. B., Dunn L., Ashton G., 2022, *MNRAS*, 511, 3304
- Newville M., Stensitzki T., Allen D. B., Ingargiola A., 2014, *LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python*, Zenodo, doi:10.5281/zenodo.11813
- Olney S., 2021, *Astron. Telegram*, 14808, 1
- Onuchukwu C. C., Legahara E., 2024, *Ap&SS*, 369, 51
- Palfreyman J., Dickey J. M., Hotan A., Ellingsen S., van Straten W., 2018, *Nature*, 556, 219
- Parthasarathy A. et al., 2019, *MNRAS*, 489, 3810
- Parthasarathy A. et al., 2020, *MNRAS*, 494, 2012
- Peralta C., Melatos A., Giacobello M., Ooi A., 2006, *ApJ*, 651, 1079
- Pons J. A., Viganò D., Geppert U., 2012, *A&A*, 547, A9
- Rabiner L., 1989, *Proc. IEEE*, 77, 257
- Reardon D. J. et al., 2023, *ApJ*, 951, L6
- Riles K., 2023, *Living Rev. Relativ.*, 26, 3
- Ruderman M., 1976, *ApJ*, 203, 213
- Sarkissian J. M., Reynolds J. E., Hobbs G., Harvey-Smith L., 2017, *Publ. Astron. Soc. Aust.*, 34, e027
- Sarkissian J., Hobbs G., Reynolds J., Palfreyman J., Olney S., 2019, *Astron. Telegram*, 12466, 1
- Shannon R. M., Cordes J. M., 2010, *ApJ*, 725, 1607
- Shannon R. M. et al., 2014, *MNRAS*, 443, 1463
- Shannon R. M., Lentati L. T., Kerr M., Johnston S., Hobbs G., Manchester R. N., 2016, *MNRAS*, 459, 3104
- Shaw B., Mickaliger M. B., Stappers B. W., Lyne A. G., Keith M. J., Weltevrede P., Basu A., 2022, *Astron. Telegram*, 15622, 1
- Singha J., Basu A., Krishnakumar M. A., Joshi B. C., Arumugam P., 2021a, *MNRAS*, 505, 5488
- Singha J., Joshi B. C., Arumugam P., Bandyopadhyay D., 2021b, *Astron. Telegram*, 14812, 1
- Sosa-Fiscella V. et al., 2021, *Astron. Telegram*, 14806, 1
- Speagle J. S., 2020, *MNRAS*, 493, 3132
- van Straten W., Bailes M., 2011, *Publ. Astron. Soc. Aust.*, 28, 1
- Taylor Joseph H. J., 1991, *IEEE Proceedings*, 79, 1054
- Urama J. O., Link B., Weisberg J. M., 2006, *MNRAS*, 370, L76
- Vallisneri M., 2020, *Astrophysics Source Code Library*, record ascl:2002.017
- Vargas A. F., Melatos A., 2023, *MNRAS*, 522, 4880
- Wong T., Backer D. C., Lyne A. G., 2001, *ApJ*, 548, 447
- Xu H. et al., 2021, *Astron. Telegram*, 14642, 1
- Yu M., Liu Q. J., 2017, *MNRAS*, 468, 3031
- Yu M. et al., 2013, *MNRAS*, 429, 688
- Yuan J. P., Wang N., Manchester R. N., Liu Z. Y., 2010, *MNRAS*, 404, 289
- Zubieta E. et al., 2022a, *Astron. Telegram*, 15638, 1
- Zubieta E., Furlan S. B. A., Palacio S. d., Garcia F., Gancio G., Lousto C. O., Combi J. A., Combi L., 2022b, *Astron. Telegram*, 15838, 1
- Zubieta E. et al., 2023, *MNRAS*, 521, 4504
- Zubieta E., Garcia F., del Palacio S., Araujo Furlan S. B., Gancio G., Lousto C. O., Combi J. A., Espinoza C. M., 2024, *A&A*, 689, A191

APPENDIX A: FLUX DENSITY CALIBRATION AND TELESCOPE SENSITIVITY FOR THE SURVEY PERIOD

We monitor the sensitivity of the telescope over the course of the timing programme using 12 bright ‘calibrator’ pulsars which exhibit relatively steady emission properties, in particular a consistent flux density from observation to observation. We use these calibrator pulsars because the NS arm upgrade did not include e.g. noise diodes, which can be used to measure variations in the sensitivity of radio telescopes (e.g. Maron et al. 2013; Jankowski et al. 2018; Kumamoto et al. 2021). The modified radiometer equation (Keane 2010) gives the proportionality between the flux density of the pulsar and the observed S/N in terms of a number of properties of the telescope and the pulsar:

$$S = \frac{T_{\text{sys}} + T_{\text{sky}}}{G \sqrt{n_p \Delta f t_{\text{int}}}} \sqrt{\frac{\delta}{1 - \delta}} (S/N), \quad (\text{A1})$$

where S is the flux density, G is the telescope gain in units of K Jy^{-1} , T_{sys} is the system temperature in K, T_{sky} is the sky temperature in K, n_p is the number of polarizations, Δf is the observing bandwidth, t_{int} is the integration time, and δ is the duty cycle of the pulsar. Thus tracking S/N for these pulsars with approximately constant S gives a proxy for the variation in the SEFD, $G^{-1} T_{\text{sys}}$, over time, as all other parameters are either constant or well-known for each observation.

The sensitivity curve we produce is quoted in relative terms – we calculate a correction factor which adjusts the measured S/N values so that they are comparable to the S/N values obtained during a period of high sensitivity between MJDs 59598 and 59652. For each calibrator pulsar, we find the mean S/N value within this calibration window, and then compute the ratio between this mean S/N and the S/N at each observation. This gives a point estimate of the telescope sensitivity, normalized against the window. All S/N values are scaled to a reference observing time of 300 s by multiplying by a factor of $\sqrt{300 \text{ s}/t_{\text{int}}}$. The final sensitivity curve is then calculated by taking the running mean over windows of 15 observations (encompassing all calibrator pulsars), and is shown in Fig. A1.

We use the telescope sensitivity curve along with a priori knowledge of the flux densities of the sample of 12 pulsars to construct an approximate flux density scale for the pulsars in our sample. To do this, we use estimates of the flux densities at 843 MHz (Jankowski et al. 2019; Johnston, private communication) for the




Figure A1. Sensitivity of the UTMOST-NS array as a function of time, normalized against the average sensitivity during the observations taken between MJD 59598 and MJD 59652, indicated by the dashed vertical lines.

same twelve calibrator pulsars used to construct the sensitivity curve. The gain of the telescope is estimated to be 0.25 K Jy^{-1} , and the system temperature approximately 170 K based on the S/N values observed from the calibrator pulsars, for an overall SEFD of 680 Jy (Mandlik et al. 2024). Given these values, along with (i) an estimate of the sky temperature for each pulsar derived from the temperature maps of Haslam et al. (1982) scaled to our observing frequency and (ii) the telescope sensitivity curve, we estimate the flux density from UTMOST-NS data for every observation of each calibrator pulsar via equation (A1), where S/N includes the correction factor derived from the telescope sensitivity curve. With the estimated system parameters quoted above, we find that the flux density estimates are in fair agreement with the Jankowski et al. (2019) values, which were calculated with reference to absolute flux density calibrated observations taken at the Parkes radio observatory (Jankowski et al. 2018). However, the UTMOST-NS flux densities tend to be somewhat lower than the Jankowski et al. (2019) values, and an overall correction factor of 1.07 is applied to the estimated SEFD to align the two sets of values, so that the mean ratio between the estimated flux densities from UTMOST-NS observations and the Jankowski et al. (2019) values is unity.

We thus obtain a flux density scale for all the pulsars in our sample, using the sensitivity curve of Fig. A1 along with the telescope parameters and overall correction factor quoted in the previous paragraph. We do not use the derived flux densities further in this work, however they are available as one of the data products through the Pulsar Data Portal⁸ (Bailes et al., in preparation). We caution that these values are indicative only, and should be treated as accurate to within a factor of approximately two.

APPENDIX B: INFLUENCE OF TIMING NOISE ON THE MEASURED $\dot{\nu}$ FOR $|\beta - 4| < 1$

This appendix presents the calculation of $\delta\dot{\nu}$ when the spectral index of the timing noise PSD, β , satisfies $|\beta - 4| < 1$. The analysis is much the same as the $\beta \approx 6$ case discussed in Section 5.1, but we assume a Langevin equation of the form in equation (21) rather than (14) – i.e. there is a random walk in ν rather than $\dot{\nu}$. The strength of the white noise term $\xi(t)$ is parametrized by σ_ν , and the phase

residual PSD is

$$P_\phi(f) = 2\sigma_\nu^2(2\pi f)^{-4}. \quad (\text{B1})$$

Again we must take into account the difference in units by multiplying the ENTERPRISE PSD by ν^2 , and account for the mismatch in spectral indices by constructing a $P^{\beta=4}(f)$ with $\beta = 4$ and A_{red} adjusted to match σ_{RN} values. The amplitude of this adjusted PSD is given by

$$A_{\text{red}}^{\beta=4} = \sqrt{\frac{3}{\beta - 1}} A_{\text{red}} \left(\frac{f_{\text{low}}}{f_{\text{yr}}} \right)^{-(\beta-4)/2}. \quad (\text{B2})$$

Then, the value of σ_ν can again be estimated by equating $P_\phi(f)$ and $\nu^2 P^{\beta=4}(f)$:

$$\sigma_\nu = \sqrt{\frac{f_{\text{yr}} \nu^2 (2\pi)^4}{24\pi^2}} A_{\text{red}}^{\beta=4}. \quad (\text{B3})$$

The dispersion in ν over a time span T is $\delta\nu = \sigma_\nu T^{1/2}$, and the corresponding dispersion in $\dot{\nu}$ is $\delta\dot{\nu} = 2\delta\nu/T^2$. This leads to

$$\delta\dot{\nu} \sim 2\sqrt{\frac{f_{\text{yr}} \nu^2 (2\pi)^4}{24\pi^2}} A_{\text{red}}^{\beta=4} T^{-3/2}, \quad (\text{B4})$$

as in equation (22).

APPENDIX C: HMM SETUP FOR OFFLINE GLITCH SEARCHES

To completely specify the HMM used for a given pulsar in the offline analysis (Section 6.4), several parameters must be chosen, including the choice of timing noise model and the parameter controlling the strength of that timing noise, as well as the parameters controlling the DOI, which is the grid of $(\nu, \dot{\nu})$ values allowed in the HMM. Note that here $(\nu, \dot{\nu})$ specifies a deviation away from a fixed timing model (with parameters indicated by a subscript LS in the main text), and so for this application, the DOI is always symmetric about zero in both ν and $\dot{\nu}$. For convenience, we introduce the uncertainty in ν , denoted $\delta\nu$, and similarly for $\dot{\nu}$. These are the uncertainties due to the ToA measurement error – the uncertainty in ν across a ToA gap of length z_k with ToA measurement errors σ_k and σ_{k+1} (in units of cycles) is given by

$$\delta\nu_k = z_k^{-1}(\sigma_k^2 + \sigma_{k+1}^2)^{1/2}, \quad (\text{C1})$$

and similarly for $\dot{\nu}$ the uncertainty is given by

$$\delta\dot{\nu}_k = z_k^{-2}(\sigma_k^2 + \sigma_{k+1}^2)^{1/2}. \quad (\text{C2})$$

The characteristic uncertainties for a given pulsar, $\delta\nu$ and $\delta\dot{\nu}$ are taken to be the median values of $\delta\nu_k$ and $\delta\dot{\nu}_k$ over all ToA gaps. Finally, we denote that median ToA gap length for a given pulsar by \bar{z} .

As noted by Dunn et al. (2023a), the nature of the HMM means that a higher cadence of observations is not always advantageous if the aim is to detect the smallest glitches in the data. Smaller z reduces the dispersion in ν and $\dot{\nu}$ during each ToA gap, but it also increases the uncertainty in ν and $\dot{\nu}$ during the gap; see equations (C1) and (C2). It therefore makes sense in some cases to remove a subset of the data before analysis, in order to minimize

$$f(\bar{z}) = \sqrt{\delta\nu^2 + (\sigma\bar{z}^\alpha)^2}, \quad (\text{C3})$$

where $f(\bar{z})$ is a measure of the average spread in ν across ToA gaps in a given data set, including the contribution from both the uncertainty due to ToA error bars ($\delta\nu$) and the timing noise model ($\sigma\bar{z}^\alpha$). The value of the exponent α depends on the timing noise model chosen;

⁸<https://pulsars.org.au>

for RWF0, we have $\alpha = 1/2$ and for RWF1 we have $\alpha = 3/2$. For each data set, we find what \bar{z} minimizes $f(\bar{z})$, and trim the data set to remove gaps smaller than this value. We impose an upper bound on the adopted optimal \bar{z} of 30 d. When no timing noise is indicated by the ENTERPRISE analysis we always take the optimal \bar{z} to be 30 d.

When choosing the HMM parameters, there are three cases to cover: the case where timing noise is not detected at a significant level (i.e. $\sigma_{\text{RN}} < 3\sigma_{\text{ToA}}$, see Section 4), the case where timing noise is significant with $2 < \beta < 5$,⁹ and the case where timing noise is significant with $\beta > 5$.

C1 No timing noise

In the case where timing noise is not detected at a significant level, we select the RWF1 timing noise model in the HMM, and set the timing noise strength parameter $\sigma_{\dot{\nu}}$ based on $\delta\nu$. Over a gap of length z , the dispersion in ν due to the random walk under RWF1 is $\sigma_{\dot{\nu}}z^{3/2}$, and so

$$\sigma_{\dot{\nu}} = \delta\nu\bar{z}^{-3/2}, \quad (\text{C4})$$

so that over the median ToA gap length \bar{z} the dispersion in ν is equal to the typical uncertainty in ν due to the ToA measurement uncertainty. The ν bounds on the DOI are taken to be $\nu_{\pm} = 10\delta\nu$, and the bin sizing to be $\epsilon_{\nu} = \delta\nu/2$. The $\dot{\nu}$ bounds are taken to be $\dot{\nu}_{\pm} = 2\nu_{\pm}/T_{\text{span}}$, and the bin spacing is $\epsilon_{\dot{\nu}} = 2\dot{\nu}_{\pm}/11$. The precise factors involved here are arbitrary but conservative – given that these pulsars do not exhibit detectable levels of timing noise, we can afford to sample ν and $\dot{\nu}$ finely, as we do not have to accommodate a wide range in either parameter.

C2 Timing noise with $2 < \beta < 5$

In the case where a model including timing noise with $2 < \beta < 5$ is preferred, we select RWF0, which produces a PSD with $\beta = 4$.

⁹The case $\beta < 2$ is excluded by our adopted priors; see Table 1.

Equation (B3) gives a recipe for setting σ_{ν} , the parameter which controls the strength of the random walk in the HMM, according to the value of A_{red} returned by ENTERPRISE. Under the RWF0 model, the dispersion in ν over a gap of length z is $\sigma_{\nu}z^{1/2}$. We take the DOI bounds on ν to be $\nu_{\pm} = \sigma_{\nu}T_{\text{span}}^{1/2}$. Similarly the DOI bounds on $\dot{\nu}$ are $\dot{\nu}_{\pm} = \sigma_{\dot{\nu}}T_{\text{span}}^{-1/2}$. The bin spacing in ν is taken to be $\epsilon_{\nu} = \max(\sigma_{\nu}\bar{z}^{1/2}/2, \delta\nu/2)$. That is, the ν bins are spaced by either the typical dispersion in ν over a ToA gap due to the random walk, or by the typical uncertainty in ν due to ToA measurement error, whichever is larger. The bin spacing in $\dot{\nu}$ is taken to be $\epsilon_{\dot{\nu}} = \min(2\dot{\nu}_{\pm}/3, \delta\dot{\nu}/2)$. We aim to space the $\dot{\nu}$ bins according to the typical uncertainty in $\dot{\nu}$ due to ToA uncertainty, but in some cases where $\dot{\nu}_{\pm}$ is particularly small this leads to only a single $\dot{\nu}$ bin. We find that the HMM produces large numbers of clear false alarms in this case, so we enforce a minimum of at least three $\dot{\nu}$ bins.

C3 Timing noise with $\beta > 5$

The final case is $\beta > 5$, and we select the RWF1 model for this case, which produces a PSD with $\beta = 6$. As with the RWF0 model discussed in Appendix C2, we have a recipe [equation (19)] for setting the strength of the random walk, in this case denoted $\sigma_{\dot{\nu}}$, based on the value of A_{red} returned by ENTERPRISE. The DOI bounds on ν and $\dot{\nu}$ are set to $\nu_{\pm} = \sigma T_{\text{span}}^{3/2}$ and $\dot{\nu}_{\pm} = \sigma T_{\text{span}}^{1/2}$, respectively. The bin spacing in $\dot{\nu}$ is taken to be $\epsilon_{\dot{\nu}} = \max(\sigma\bar{z}^{1/2}/2, \delta\dot{\nu}/2)$. As for the case of ν spacing in Appendix C2, the $\dot{\nu}$ bins are spaced either according to the typical dispersion in $\dot{\nu}$ over a ToA gap or by the uncertainty in $\dot{\nu}$ due to ToA uncertainty, whichever is larger. In some cases, the spacing in $\dot{\nu}$ and the initially chosen value of $\sigma_{\dot{\nu}}$ are such that transitions to neighbouring $\dot{\nu}$ bins are suppressed, and we find that this causes the HMM to produce large numbers of clear false alarms. A similar issue arises with $\dot{\nu}$ spacing in Appendix C2. We include an adjustment to $\sigma_{\dot{\nu}}$ at this point, taking $\sigma_{\dot{\nu}} = \max(\sigma_{\dot{\nu}}, \epsilon_{\dot{\nu}}\bar{z}^{-1/2})$, so that $\sigma_{\dot{\nu}}$ is large enough to allow the hidden state to transition between neighbouring $\dot{\nu}$ bins. Finally, we take the bin spacing in ν as $\epsilon_{\nu} = \max(\sigma\bar{z}^{3/2}, \delta\nu/2)$.

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