

Observation of discontinuities in the periodic modulation of PSR B1828–11

Adriana Dias¹,¹★ Gregory Ashton,¹ Julianna Ostrovska,¹ David Ian Jones² and Michael Keith³

¹*Physics Department, Royal Holloway, University of London, Egham Hill, Egham TW20 0EX, UK*

²*Mathematical Sciences, University of Southampton, Southampton SO17 1BJ, UK*

³*Department of Physics and Astronomy, Jodrell Bank Centre for Astrophysics, The University of Manchester, Manchester M13 9PL, UK*

Accepted 2025 March 19. Received 2025 March 19; in original form 2025 January 20

ABSTRACT

PSR B1828–11 is a radio pulsar that undergoes periodic modulations (~ 500 d) of its spin-down rate and beamwidth, providing a valuable opportunity to understand the rotational dynamics of neutron stars. The periodic modulations have previously been attributed to planetary companion(s), precession, or magnetospheric effects and have several interesting features: they persist over 10 cycles, there are at least two harmonically related components, and the period is decreasing at a rate of about 5 d per cycle. PSR B1828–11 also experienced a glitch, a sudden increase in its rotation frequency, at 55040.9 MJD. By studying the interaction of the periodic modulations with the glitch, we seek to find evidence to distinguish explanations of the periodic modulation. Using a phenomenological model, we analyse a data set from Jodrell Bank Observatory, providing the longest and highest resolution measurements of the pulsar’s spin-down rate data. Our phenomenological model consists of step changes in the amplitude, modulation frequency, and phase of the long-term periodic modulation and the usual spin-down glitch behaviour. We find clear evidence with a (natural-log) Bayes factor of 1486 to support that not only is there a change to these three separate parameters but that the shifts occur before the glitch. We also present model-independent evidence which demonstrates visually how and when the modulation period and amplitude change. Discontinuities in the modulation period are difficult to explain if a planetary companion sources the periodic modulations, but we conclude with a discussion on the insights into precession and magnetospheric switching.

Key words: stars: neutron – pulsars: general – pulsars: individual: PSR B1828–11.

1 INTRODUCTION

Pulsars provide a unique astrophysical laboratory to probe physics at the extreme. One avenue to better understand pulsars is through the investigation of pulse timing, which may reveal insights into the properties of the magnetosphere (which emits the observed radiation) or the interior of the neutron star itself. In this work, we study data on the spin-down rate of PSR B1828–11 (i.e. the time derivative of the pulsation frequency), performing a phenomenological model fit to study features in a new high-resolution data set recorded at Jodrell Bank Observatory (JBO). This pulsar exhibits several interesting and related phenomena: the timing properties are periodically modulated with a time-scale of ~ 500 d and display a characteristic double-harmonic-sinusoid structure. Meanwhile, the pulse shape rapidly switches between two distinct states, and the proportion of time spent in each state is also modulated and correlated with the timing variations. Finally, the modulation period decreases with time, and the star has undergone a glitch – a sudden spin-up event. This rich mixture of observations requires a unified explanation. Three primary model interpretations have been proposed: the presence of a planet or system of planets orbiting the pulsar, free precession, and

magnetospheric switching. The ultimate goal of this work is to utilize the new data to constrain these models.

2 PREVIOUS STUDIES OF PSR B1828–11

Bailes et al. (1993) reported the first observations of PSR B1828–11 and hypothesized a planetary explanation, noting that a system of at least two planets would be required to explain the two harmonics observed in the timing properties of the star. However, in Stairs, Lyne & Shemar (2000), an extended data set was analysed, covering several cycles and simultaneously analysing timing properties and the pulse shape (via an averaged shape parameter $\langle S \rangle$), finding strong correlations between the two. Based on this observation, Stairs et al. (2000) rejected the planetary explanation since it would require the planet, orbiting at about 1 au, to interact with the magnetosphere that is at most a few thousand kilometres. Nevertheless, recent work by Liu, Yue & Xu (2007) studied a quark planetary model and separately Nițu et al. (2022) conducted a search for planetary companions around 800 pulsars, finding that PSR B1828–11 could, in principle, be explained with two planetary companions (though they conclude that intrinsic spin variation is a better-supported explanation).

Instead, Stairs et al. (2000) proposed free precession as the cause of the periodic modulation. They postulated that the periodicity of

* E-mail: adriana.dias@rhul.ac.uk




Figure 1. Comparison between Lyne et al. (2010) (in the foreground) and Keith & Niu (2023b) (in the background, used in this work) data sets of spin-down rate, with respect to time in MJD. The dashed vertical line highlights the glitch time for PSR B1828–11.

PSR B1828–11, with harmonically related sinusoids with periods of ~ 1000 , 500, and 250 d, was resultant from precession of the spin axis, caused by the misalignment of the angular momentum and symmetry axis of the star and assuming the star to be non-spherical. This work was followed by physical models proposed by Jones & Andersson (2001) and Link & Epstein (2001), where the authors explored how the variations in the pulse shape and timing of PSR B1828–11 could be explained by free precession of the star’s crust causing variations in the magnetic dipole torque angle. They found the observations could be explained by a star precessing with a period of ~ 500 d and a wobble angle of $\sim 3^\circ$, assuming that the magnetic dipole is nearly orthogonal to the star’s symmetry axis. This configuration is somewhat special as it means that the dipole cuts through the equator four times per precession period, producing the characteristic double-harmonic-sinusoid observations (see Fig. 1). Moreover, Link & Epstein (2001) fitted the model to the data and found that an hourglass-type beam geometry was required to explain the observed $\langle S \rangle$ data. Further advances of precession include a tri-axial body with core and blob beam geometry (Akgün, Link & Wasserman 2006) and the development of a time-varying magnetic field (Rezania 2003).

Following further observations, the free precession interpretation was challenged in Lyne et al. (2010). Most notably, they highlight that the time-averaging baseline required to measure the spin-down rate (used in the beam-shape parameter of Stairs et al. 2000) will obscure behaviour happening on faster time-scales. Following the contemporaneous identification of rapid magnetospheric switching phenomena (see e.g. the extreme case of PSR B1931+24 Kramer et al. 2006, where the pulsar switches on and off with correlated changes in its spin-down rate), the authors proposed that the spin-down and beamwidth variations of PSR B1828–11 could similarly be explained by a model in which the magnetosphere switches between states in a quasi-periodic fashion, but that the probability of being in one state or the other varies on the modulation time-scale. This suggestion is based on the identification (Stairs et al. 2000, 2003) that the pulsar exhibits distinct narrower and wider profiles. They explained that inferred parameters such as the spin-down rate and shape parameter $\langle S \rangle$, which use a multiday baseline, average over behaviour on shorter time-scales, revealing the slow time-varying

probability between states but obscuring the rapid switching. To evidence this, they introduce a new pulse shape parameter, w_{10} , that could be calculated on individual observations to avoid the longer 100-d baseline required to measure the spin-down. A further follow-up study of PSR B1828–11 in Stairs et al. (2019) used additional high-resolution observations from the Parkes and Green Bank Telescopes, which enabled a detailed study of the pulse-to-pulse behaviour. They confirmed that there are only two distinct pulse shapes. By correlating the ratio of the time spent in each state with the modulation phase, they further validate the model proposed by Lyne et al. (2010).

A generative model of the switching process was developed in Perera et al. (2015) and applied to the PSR B0919+06, which shows a similar pattern of long-term behaviour to PSR B1828–11. To explain the characteristic double-harmonic-sinusoid present in the spin-down rate of B0919+06 with a two-state magnetospheric switching model, Perera et al. (2015) proposed a *four-phase model* in which the pulsar switches between the two states twice per cycle. The characteristic second (lower) peak arises because the time spent in the state is shorter than the time-averaging window used to generate the spin-down data. The authors of Shaw et al. (2022) also report a similar behaviour in PSR B0740–28, whose profile exhibits two distinct shapes. Taking this model, Ashton, Jones & Prix (2016) compared the precession and switching hypotheses for PSR B1828–11, analysing the spin-down and w_{10} pulse shape data from Lyne et al. (2010). They augmented the standard precession model with a variable braking index and included a flexible beam profile. Meanwhile, the four-state switching model proposed by Perera et al. (2015) was applied, additionally modelling the time-averaging process to predict the spin-down and connecting each state with a separate beamwidth w_{10} . Ashton et al. (2016) concluded that, based on the models and data under consideration, precession was the favoured explanation.

However, the more recent study of Stairs et al. (2019) points out that the precession beamwidth model applied in Ashton et al. (2016) is at odds with the observation. Specifically, w_{10} varies slowly due to changes in the line-of-sight view of the radio emission, while the data demonstrate that it, in fact, varies rapidly between pulses. Nevertheless, while the pulse shape model of Ashton et al. (2016) is mistaken, the precessional explanation of the spin-down is the

more parsimonious: it both provides a natural clock and avoids the complicated four-phase model required to explain the double-harmonic-sinusoid. Moreover, as pointed out by Jones (2012), it remains plausible that precession is the clock driving the long-time-scale variability. For example the spin-down could remain a product of the effects of precession, while the time-varying wobble of the star could be responsible for driving the unstable magnetosphere to switch between quasi-stable states.

We also point out in reviewing Ashton et al. (2016) that the Perera et al. (2015) four-phase switching model also seems to be at odds with our new observations. Specifically, in this model, the pulsar switches rapidly between two distinct spin-down rates but switches twice per cycle. To produce the double-peaked spin-down rate, the time spent in one of the states must be shorter than the time-averaging baseline. It therefore follows that the secondary peak ‘height’ is a function of the time-averaging baseline. If the baseline is sufficiently short, there will be a double peak, but the heights will be equal; only the duration spent in each state will differ.

In Fig. 1, we can compare the data produced using the 100 d baseline from Lyne et al. (2010) with the higher resolution data obtained generated by Keith & Ntũ (2023b) using a Fourier-basis Gaussian process regression (GPR) (described later in Section 3). Notably, we do not see a variation in the height of the second peak. While the methods are not directly comparable, the average time between time of arrivals (TOAs) in the data analysed by Shaw et al. (2022) was 7 d; since in Ashton et al. (2016), it was shown that the duration spent in the short-duration state was approximately 10–20 d, we would therefore expect the Keith & Ntũ (2023b) inferred spin-down measurements to be more sensitive to the step-changes if the pulsar switches suddenly and semipermanently between states as in the Perera et al. (2015) model. However, this is not the case. Therefore, this observation is inconsistent with the four-phase switching model and suggests that whatever mechanism drives the spin-down variations smoothly varies between the minima and maxima (as previously argued and demonstrated by Stairs et al. 2019). To account for this observation, the four-phase switching model could be modified. Minimally, one could introduce three distinct states, though this then causes inconsistencies with the observation of the beamwidth, which itself does not show any evidence of a third state.

The complexity of B1828–11 became more interesting when Ashton et al. (2017) discovered that the modulation period present in the spin-down data is itself getting shorter, losing about 1 d per 100 d, and identified from the Jodrell bank glitch catalogue (Basu et al. 2021) that the pulsar also experienced a glitch, a sudden increase in the rotation frequency, at 55040.9 Modified Julian Day (MJD) (coinciding with the end of the data set provided by Lyne et al. 2010). Leading models of glitches suggest they provide evidence for a superfluid component in the core of the star. However, such a superfluid component is incompatible with precession (Shaham 1977; Jones & Andersson 2001; Haskell & Jones 2024) since the pinning of the superfluid would result in a free precession period much shorter than the observed modulation period of ~ 500 d, and may also be expected to be rapidly damped.

The implications of this were discussed in Jones, Ashton & Prix (2017), where several models tried to tie together the decreasing modulation period with the glitch, making predictions for the subsequent behaviour. In the main, these predicted that the glitch should produce changes in the modulation period. However, subsequent analyses (Brook et al. 2016; Stairs et al. 2019; Shaw et al. 2022) have demonstrated that the modulation of the timing properties on a ~ 500 d time-scale continues after the glitch. However, to date, no

quantitative study has been performed to determine if there are any step changes associated with the glitch.

Very recently, Lower et al. (2025,) made a study of radio emission variability in a sample of 259 pulsars, making two findings of potential relevance here. First, they found that variations in both spin-down rate and pulsar profile shape are more common than previously thought. Secondly, by looking at the set of 45 pulsars that exhibit quasi-periodic variations in their spin-down date, they found that the modulation period of the variations was approximately independent of the spin period, a result not expected on the basis of several free precession models described in Jones (2012). This last point makes the free precession interpretation of quasi-periodic timing variability less attractive, at least as a common explanation for all such variable pulsars.

In any case, it remains unclear what mechanism is responsible for the long-term behaviour of PSR B1828–11. While we can agree that the magnetosphere switches rapidly between two states and that this varies coherently on a 500-hundred-day cycle with variations in the spin-down, we do not yet know ‘what sets the clock of this cycle’? ‘If it is switching between just two states, why is the spin-down smoothly varying’? and ‘Why is the modulation period decreasing, and are there any changes related to the glitch’?

To answer these questions, we revisit the analysis of PSR B1828–11 using the new high-resolution spin-down data (Keith & Ntũ 2023b). And, to avoid preconditioning our interpretation with a physical model, we apply a phenomenological model to capture the salient features that may be present in the spin-down rate of the pulsar. We will model the spin-down rate data for this pulsar and ascertain whether any step changes occur around the glitch that changes its spin-down rate or modulation. To consider several possibilities, we developed three models to describe the spin-down rate of this pulsar: a model which assumes that a glitch occurred and that there are changes to the periodic modulations; a model which assumes that there is no glitch nor changes to the periodic modulation; and another one which assumes there is a glitch but no changes to periodic modulation. We obtained and compared the natural-log evidence for these three models to understand which one fits the data more appropriately.

The paper is structured as follows. We first introduce the data set and methodology in Sections 3 and 4 before describing the models and the fits to the data in Section 5. Then, in Section 6, we study the time-period behaviour of the pulsar and compare this with the features extracted from the model. Finally, we conclude with a discussion and outlook in Sections 7 and 8, respectively.

3 DATA

In this work, we will analyse the open spin-down data published in Keith & Ntũ (2023b) (and available from Keith & Ntũ 2023a), which was derived using a Fourier-basis GPR on the raw data in Shaw et al. (2022). Observations were conducted using the 76 m Lovell telescope and were supplemented with data from the 25 m ‘Mark-II’ telescope, both located at JBO (Lovell 1957). Data collected before 2009 were centred at 1400 MHz and recorded using a 32 MHz filterbank. After 2009, data collection shifted to being centred at 1520 MHz and recorded with a 384 MHz filterbank. Detailed information on data acquisition settings can be found in Shaw et al. (2022). To transform the acquired data into the spin-down rate analysed in this paper, Shaw et al. (2022) generated a single pulse profile for each observation epoch by summing the data across all frequency channels. The TOA is obtained by comparing this integrated pulse profile with a high signal-to-noise profile representing the observed profile’s expected shape. The TOAs are then fitted with a timing

model (Hobbs, Edwards & Manchester 2006), then subtracting this model from the data results in a timing residual. Finally, the timing residual is fitted using Fourier-basis GPR and from this, the second derivative of the spin-down $\dot{\nu}$ is extracted (see Keith & Ntũ 2023b for further details).

Previous studies on PSR B1828–11, i.e. in Ashton et al. (2017), utilized a smaller data set spanning 5280 d between 49710 and 54980 MJD, which ended before the glitch occurred. This data set had 755 TOAs, and the spin-down rate was obtained by applying a timing model to a sliding window of duration 100 d over the data (Lyne et al. 2010). In contrast, the Keith & Ntũ (2023b) data set used in this paper spans 8615 d between 49202 and 57817 MJD, encompasses the glitch event, and has better resolution. Fig. 1 illustrates the differences between these data sets, with the old data set shown in orange, the newer data set in blue and the glitch time marked by a vertical black dotted line.

4 DATA ANALYSIS METHODOLOGY

This section provides a brief overview of the Bayesian methodology we use to analyse the data under a set of phenomenological models. (For a general introduction, see e.g. Carlin et al. 2013).

Bayes theorem aims to solve the inverse problem: what can be learnt about model M and its associated parameters ϑ , based on data d ? This can be described by equation (1):

$$p(\vartheta|d, M) = \frac{\mathcal{L}(d|\vartheta, M)\pi(\vartheta|M)}{\mathcal{Z}(d|M)}, \quad (1)$$

where $p(\vartheta|d, M)$ is the posterior probability distribution of the parameters ϑ given the data and the model; $\mathcal{L}(d|\vartheta, M)$ is the likelihood function of the data, given the parameters and the model; $\pi(\vartheta|M)$ is the prior probability distribution, associated with the set of model parameters; and $\mathcal{Z}(d|M)$ is the evidence for the data, given the model, and can be calculated from $\mathcal{Z}(d|M) = \int_{\vartheta} \mathcal{L}(d|\vartheta, M)\pi(\vartheta|M)d\vartheta$.

We will use the Bilby Bayesian inference library (Ashton et al. 2019) to draw samples from the posterior probability density and estimate the evidence using the nested sampling algorithm (Skilling 2004), specifically, the *dynesty* sampler (Speagle 2020; Koposov et al. 2024). Nested sampling enables efficient exploration of the multimodal and higher dimensional space we will explore, producing a set of samples approximating the posterior $p(\vartheta|d, M)$ and an estimate of the evidence $\mathcal{Z}(d|M)$ which we will use for model comparisons.

In contrast to previous works that used Bayesian analyses, throughout this work, we will use ‘slab-and-spike’ priors (Malsiner-Walli & Wagner 2016). These comprise a *slab*, usually a standard prior distribution, such as a uniform or a normal distribution prior, and a Dirac *spike* at a fixed location. We use these in our phenomenological model as a means to marginalize over the model dimensionality without requiring the implementation of a transdimensional sampler (Green 2003). Consider a polynomial of degree N with coefficients a_i with $i \in [0, N]$ as a generic example. A naive analysis may apply a Bayesian analysis to each degree, treating each as a separate ‘model’; a transdimensional sampler improves on this by including N as a model parameter, enabling automatic marginalization over the model size. However, implementation is often domain-specific (though see Tong et al. 2025). Instead, slab-spike priors can be used with regular stochastic samplers when the models are nested (e.g. in the polynomial case, a model of degree N is equivalent to a model with degree $N - 1$ with the parameter a_N fixed to zero). By placing the spike at the point that recovers the simpler model (e.g. $a_i = 0$), higher dimensional models can be explored with the sampler finding

posteriors equal to zero for higher dimensional parameters that do not improve the fit.

5 DEFINING AND FITTING MODELS

In this section, we define three phenomenological models of the secular spin-down and periodic modulations to fit the data in Fig. 1. For all three models, the secular part encodes a standard expansion of N_s frequency derivatives, and the periodic modulations utilize a sinusoid with N_c harmonically related components. Within each sinusoidal term, the phase follows an expansion up to the N_f phase-derivate to capture the slow changes to the modulation period observed in Ashton et al. (2017). We start with the most general model, referred to as **Model: S+P**, which allows independent step changes in the secular spin-down and periodic modulation. We also explore two subsets of the **S+P** model: one which assumes that there is no glitch nor changes to the secular spin-down or the periodic modulation (**Model: no-glitch**) and another one which assumes a step change only in the secular spin-down (**Model: S**). These subset models allow us to probe the significance of changes in the periodic modulation relative to the other step changes. For each model, we discuss the theoretical reasoning first, then explain the choice of priors and, finally, the inferred posteriors.

5.1 Model: S+P

In this section, we define and apply a model in which a glitch occurs (modelled by an instantaneous change in the spin-down rate accompanied by a transient decay) and that there are also instantaneous changes to the features of the periodic modulation. We model changes to the features of the periodic modulations as step functions and allow a step change in each component separately.

To develop a full generative model, first we define $t' = t - t_0$, where t is the MJD of the observed data and t_0 is the MJD of a reference time (55372 MJD as quoted in Parthasarathy et al. 2019). We then write the spin-down rate as:

$$\begin{aligned} \dot{\nu}(t) = & \sum_{i=0}^{N_s-1} \frac{\dot{\nu}_i}{i!} \left[1 + H(t' - t_s^{\xi}) \left(\xi_i^p + \xi_i^t e^{-\frac{t' - t_s^{\xi}}{\tau_i}} \right) \right] \Delta t^i \\ & + \sum_{j=1}^{N_c} A_j \left[1 + H(t' - t_s^{\eta}) \eta_j \right] \\ & \cdot \cos \left(j\phi(t) + \Delta\phi_j \left(1 + H(t' - t_s^{\delta}) \delta_j \right) \right), \end{aligned} \quad (2)$$

where the phase is given by

$$\phi(t) = 2\pi \sum_{k=0}^{N_f-1} \frac{1}{k!} f_k \left(1 + H(t' - t_s^{\chi}) \chi_k \right) \Delta t^{k+1}. \quad (3)$$

The key components of this model are: $\dot{\nu}_i$, the i th coefficient of the spin-down expansion, $\dot{\nu}_i = \frac{d^{(i)}\dot{\nu}}{dt^{(i)}}$; A_j , the j th cosine component coefficient (amplitude); $\Delta\phi_j$, the phase-offset of the j th cosine component; and f_k , the k th derivative of the modulation frequency.

Within this model, the parameters N_s , N_c , and N_f define the maximum number of components included in the model. Ideally, we would like to marginalize over these parameters (e.g. using a transdimensional sampler). However, in practice, we will use a maximum value and then apply slab-spike priors. To determine the maximum value, we analysed the data using the **S+P** model, incrementing each maximum until no improvement in the fit was found (as quantified by the change in the Bayesian natural-log evidence). Using this approach, we selected maximum values of

$N_s = 3$, $N_c = 8$, and $N_f = 2$; the choice of upper limit on the number of frequency components is consistent with the frequency range of the Fourier-basis GPR used to generate the data (Keith & Nijū 2023b). To confirm these were sufficiently large, we then verified that the amplitude parameter of the largest component had the maximum posterior support at zero (see Table A1); in other words, the model preferred a simpler model, and our results are robust to increases in the maximum values.

To model the step changes in each component of equations (2) and (3), we utilize a Heaviside step function multiplying a dimensionless relative amplitude for the spin-down ξ^p , the amplitude η , modulation phase δ_j , and modulation frequency χ_k . There is one step-change per summation term, but all parameters in a given set have a single associated time (i.e. t_s^ξ , t_s^η , t_s^δ , and t_s^χ) such that, e.g. all spin-down changes happen at the same time. For the changes to the secular spin-down rate, following standard approaches to modelling glitches (see e.g. Lorimer & Kramer 2005), there is a permanent offset ξ^p and an exponentially decaying transient term ξ^t with associated time-scale τ . However, the transient component is only non-zero for the zeroth-order term as set by the priors discussed shortly.

The modulation component of the model follows a harmonic-sinusoid form with an amplitude A_j , harmonic phase $j\phi(t)$ (where $\phi(t)$ is the standard phase expansion), and phase offset $\Delta\phi_j$. The harmonic coefficient j multiplies the phase in the argument of the cosine but does not multiply the phase offset. This prevents degeneracies in the solution as $\Delta\phi \in [0, 2\pi]$ while still exploring the entire parameter space.

Priors. We list the complete set of priors used in Table A1. For many parameters, we use a uniform prior, choosing a suitable range to cover the expected behaviour (and check where performed to ensure the range did not arbitrarily limit the model fit). We then augment several of these with slab-spike priors emulating a transdimensional sampler.

For the glitch time parameter t_s^ξ affecting the secular spin-down, we apply a prior width ranging ± 50 d around the 55040.9 MJD based on the recorded glitch time (Basu et al. 2021). Meanwhile, for the other step-change time parameters, we sample in an offset time relative to t_s^ξ : that is we define $\Delta t_s^\alpha = t_s^\alpha - t_s^\xi$ for $\alpha \in \{\chi, \delta, \eta\}$ and then apply a uniform prior on Δt_s^χ , Δt_s^δ , and Δt_s^η from -5000 to 2000 d.

For ν_0 , we apply a wide prior ranging from the minimum to the maximum values of the observed spin-down data shown in Fig. 1, i.e. from -2.74×10^{-3} to $-2.72 \times 10^{-3} \text{ d}^{-2}$. For all higher order derivatives of ν_0 , we set a uniform prior on an arbitrary range and verify the choice of the prior range has no impact on the analysis.

For ξ_p^0 and ξ_t^0 , we set a uniform prior with a range ± 0.01 and again verify this arbitrary range is sufficiently broad. For τ , we apply a uniform prior between 0 and 500 d, ensuring the relaxation time is positive while choosing an arbitrarily large upper value.

The amplitude terms, A_j , are given prior distributions ranging from 0 to 10^{-5} ; while a negative amplitude is, in principle, physical, this would introduce degeneracy with the phase term. The modulation phase offset terms are given a uniform prior on $-\pi$ to π . For the step-change parameters, we apply a uniform prior from -1 to 1 for χ_j ; we set a uniform prior on -1 to 1 ; this enables direct interpretation of the posterior without concern about the effects of the prior. However, for η_j and δ_j , we found that with a uniform prior, the sampler failed to robustly identify the maximum-posterior mode (occasionally getting stuck in islands with lower posterior support with larger relative changes. Therefore, we instead apply a standard normal prior such that the prior maximum is zero while setting a scale for expected instantaneous changes, which

suppresses order-of-magnitude increases in the amplitude and phase term.

The phase, as seen in equation (3), contains k th derivatives of the modulation frequency. The base modulation period is estimated to be ~ 460 d, although as it is shown in Fig. 10 this modulation period varies from 489 to 435 d throughout the entire data range. Thus, we set the prior range of f_0 to a range which includes the base modulation frequency, i.e. $\frac{1}{460}$ Hz. The other f_k terms have an arbitrary factor of 10^{4+k} applied to the modulation frequency.

Results. We summarize the posterior distributions in Table A4, which contains the median \pm standard deviation values. Fig. 2(a) presents the spin-down rate data (in blue) together with the maximum posterior estimate solution of the model (in red) and an orange dashed line showing the secular component of the model alone (i.e. without the periodic modulation) from which we see the analysis identified an exponential recovery present after the glitch. Additionally, four vertically shaded 99 per cent quantile regions are shown, which relate to each of the step-change time parameters, with t_s^ξ shown in blue, t_s^η shown in yellow, t_s^δ shown in green, and t_s^χ shown in red. A detailed description of the results of the glitch step change is presented in Section 5.2.

In Fig. 2(b), we visualize the residuals obtained by subtracting the model from the data alongside the 90 per cent interval generated by sampling model draws from the posterior distribution before subtracting for the residual. We note that, while the broad fit to the data is good, the residual still displays some structure, suggesting further improvements to our phenomenological model are possible.

5.2 Interpreting the step-changes in inferred parameter for Model S+P

We find that the posterior distribution of t_s^ξ has a posterior width of ~ 6 d at the 99 per cent credible interval (see Fig. 3), with a maximum posterior value of 55049, ~ 9 d apart from the recorded glitch time of PSR B1828–11 (shown as a vertical dashed line in Fig. 2a). This difference likely arises from the fact that we are estimating the glitch time from the spin-down rate whereas the glitch time is estimated from the full phase evolution.

For the secular spin-down, we measure the spin-down and its first two derivatives with values consistent with those known in the literature. We also measure a third-order derivative that while non-zero, contained zero at 3 standard deviations. Fig. 2(a) shows that the model has recovered the step-change and transient recovery observed on the data (which has an inferred time-scale of 121 ± 14 d. We do not allow for step changes in the spin-down derivatives and there is no evidence from the residuals to suggest these are required.

For the periodic modulations, we identify eight non-zero harmonic components, a significant change relative to the two harmonic components that have been fitted to the data before (see e.g. Ashton et al. 2016; Stairs et al. 2019). The impact of these higher-order terms can be observed directly in Fig. 1: looking at the trailing edge after each of the successive maxima, we can identify in the data a short plateau; this was present in the original data set (Lyne et al. 2010), but is distinct in the newer higher-resolution data analysed in this work. In Fig. 2(a), we see the corresponding behaviour of the higher-order terms in the harmonic expansion fitting this feature (this is also present in fits of the no-glitch model as well, c.f. Fig. 8a).

From our analyses, we also identify that a shift in the modulation amplitude occurs at $t_s^\eta = 54316.34$ MJD before the glitch occurs. To visualize the posterior distributions, in Figs 4 and 5, we plot the posterior distribution for the amplitudes $A_{[1-9]}$ and relative amplitude changes $\eta_{[1-9]}$, respectively. Notably, the first component undergoes




Figure 2. (a) Shows the spin-down rate data, in blue, together with the maximum posterior estimate solution of **Model S+P**, in red, which uses the parameters with the highest posterior probability. An orange dotted line shows the spin-down rate component of the model without the modulation cosine components. The glitch time is represented by a black dashed vertical line. Four vertical shaded 99 per cent quantile regions are shown, which relate to each of the t_s step-change parameters, with t_s^ξ (the step in spin-down rate) shown in blue, t_s^η (the step-in modulation amplitude) shown in yellow, t_s^δ (the step-in modulation phase) shown in green, and t_s^χ (the step-in modulation frequency) shown in red. (b) Shows the residuals, as a line in blue, obtained by subtracting **Model S+P** from the data. The blue shaded area around the data shows the 90 per cent quantile region. Here, a dashed vertical line also indicates the glitch time.

a ~ 20 per cent decrease in amplitude while the second component increases by about the same amount. These two components are the leading order, and the impact can be seen by comparing the fit before and after the glitch in Fig. 2(a).

We find that at the highest-order component (i.e. A_9 and η_9), the posterior peak is at zero, justifying that our choice of N is sufficiently large. However, while this suggests the existence of multiple higher-order harmonics (above the two fitted in previous works), we do note that this could arise from the GPR methodology used to infer the spin-down rate (Keith & Nițu 2023b).

We also find evidence for a distinct step-change in the modulation phase at $t_s^\delta = 50622.26$ MJD. However, by eye, it is difficult to distinguish in Fig. 2(a) what feature this is fitting: there is no clear discontinuity in the phase at this time.

We recover a modulation frequency and non-zero first derivative consistent with values already reported in the literature (Ashton et al. 2017). Our model is also sensitive to a second derivative not previously explored – Fig. 6. However, the posterior distribution is consistent with zero, i.e. we do not find any evidence for a second derivative of the modulation period. We find evidence for a distinct step change in the modulation frequency at $t_s^\chi = 53615.11$ MJD, some 1434 d before the glitch time t_s^ξ (55048.92 MJD). The posterior distributions show that the modulation period and its first derivative experience fractional shifts of 3.92×10^{-04} and 8.07×10^{-01} , respectively. To visualize this behaviour, in Fig. 7, we plot the inferred modulation period as a function of time.

To test the significance of our discovery that the modulation frequency, phase, and component amplitudes change at disjoint times




Figure 3. Posterior distribution for the t_s^ξ parameter that gives the time of the glitch as found in our **Model S+P**. The 99 percent quantile region is shown and the previously reported glitch time of PSR B1828–11 is represented by the vertical dashed line.




Figure 4. Posterior probability distribution for the A_j parameters inferred from fitting **Model S+P**.




Figure 5. Posterior probability distribution for the η_j parameters inferred from fitting **Model S+P**. Note η_5 and η_6 are bimodal distributions with a mode at zero and a non-zero mode.

spanning nearly the entire data set, we repeat the analysis but restrict the prior distributions on the times of the step-changes to ± 50 d, centred on the reported glitch time. We find that the Bayes factor between the full model and this restricted analysis is decisively in support of the full model with a natural-log Bayes factor of 877.

5.3 Model subsets

To probe the relative importance of different features of the **S+P** model, we now explore two model subsets. First, a model which assumes that there is no glitch nor changes to the periodic modulation (**Model no-glitch**) and then a model which includes a step change only in the secular spin-down (**Model S**).

For **Model no-glitch**, we modify equations (2) and (3) removing the step changes leading to

$$\dot{\nu}(t) = \sum_{i=0}^{N_s-1} \frac{\dot{\nu}_i}{i!} \Delta t^i + \sum_{j=1}^{N_c} A_j \cos(j\phi(t) + \Delta\phi_j), \quad (4)$$

and

$$\phi(t) = 2\pi \sum_{k=0}^{N_f-1} \frac{1}{k!} f_k \Delta t^{k+1}. \quad (5)$$

Meanwhile, for **Model S**, which assumes a step-change at the glitch for the spin-down rate, we include a step change only in the secular part of the spin-down, i.e.:

$$\begin{aligned} \dot{\nu}(t) = & \sum_{i=0}^{N_s-1} \frac{\dot{\nu}_i}{i!} \left[1 + H(t' - t_s^\xi) \left(\xi_i^p + \xi_i^t e^{-\frac{t' - t_s^\xi}{\tau_i^\xi}} \right) \right] \Delta t^i \\ & + \sum_{j=1}^{N_c} A_j \cos(j\phi(t) + \Delta\phi_j) \end{aligned} \quad (6)$$

with

$$\phi(t) = 2\pi \sum_{k=0}^{N_f-1} \frac{1}{k!} f_k \Delta t^{k+1}. \quad (7)$$

Model results. The procedure described in Section 5.1 was applied to both model subsets, from how the priors were defined to how the posterior distributions were obtained. Tables A2 and A3 list the full set of priors, for **Model no-glitch** and **Model S**, respectively. The choice of priors was the same as the ones presented in Section 5.1 for **Model S+P**, but no glitch-related parameter priors and no modulation change-related parameter priors were included, for **Model no-glitch** and **Model S**, respectively.

These subsets were obtained with $N_s = 4$, $N_c = 9$, and $N_f = 3$ and Tables A5 and A6 show that f_2 , A_9 , and $\dot{\nu}_3$ have maximum posterior probability values consistent with 0 within 1σ , as was previously noted in **Model S+P**. **Model no-glitch** and **Model S** returned natural-log evidences of -68308.4 ± 0.2 and 68445.6 ± 0.2 , respectively, lower than what was obtained for **Model S+P** (69931.9 ± 0.2).

Figs 8(a) and 9(a) present the spin-down rate data (in blue) together with the maximum posterior estimate solution of the model (in red), which uses the parameters with the highest posterior probability, for **Model no-glitch** and **Model S**, respectively. These subset models were not able to capture the changes in the spin-down rate, in particular the transient recovery, that occurred after the glitch. This is evident in these figures but also in Figs 8(b) and 9(b), which show the residuals obtained by subtracting **Model no-glitch** and **Model S** from the data, respectively. By comparing these figures with Fig. 2(b), we can see that **Model no-glitch** and **Model S** are unable to capture the

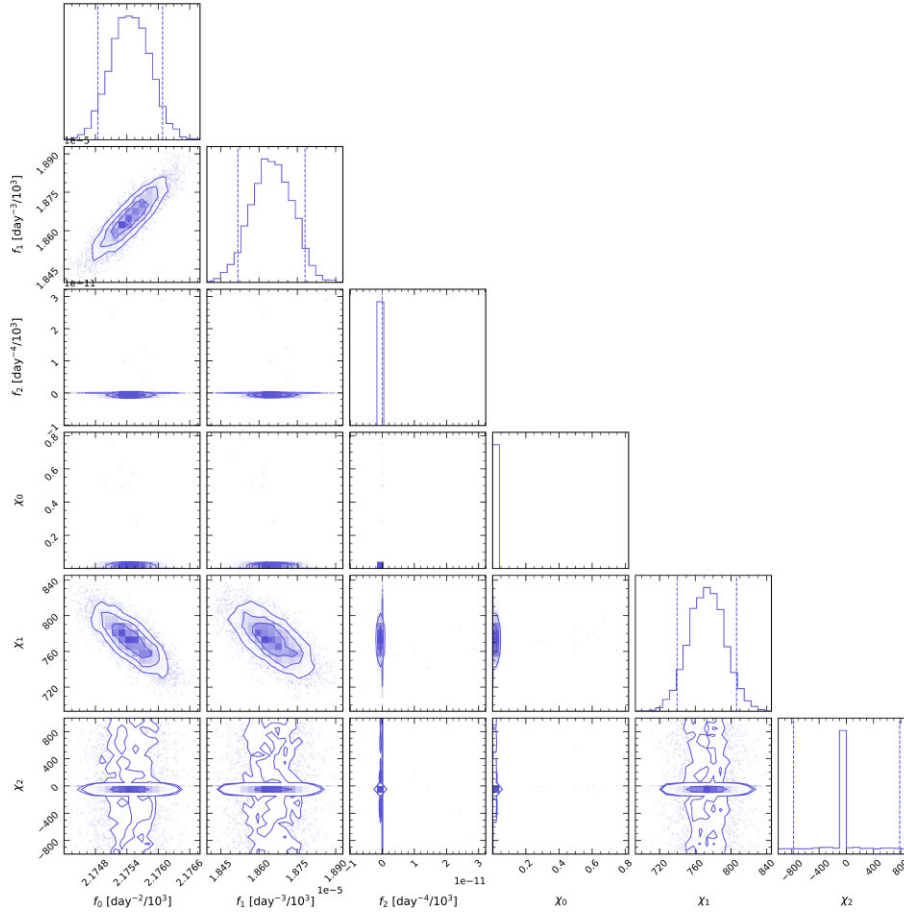


Figure 6. Posterior probability distribution for the modulation frequency terms, f_k , and χ_k , which represent their step change, for **Model S+P**.




Figure 7. Modulation period versus MJD, for **Model S+P**. The bold vertical line indicates the t_s^X glitch time parameter, and the dotted line indicates the glitch time.

changes occurring on the data after the glitch, in the region between 55000 and 56000 MJD.

Discussion. The model subsets perform poorly relative to **Model S+P** in modelling the observed spin-down rate of PSR B1828–11. This is evidenced by the residual plots, which show larger deviations from zero and clear structures indicating specific instances where

they fail, such as the transient recovery after the glitch, as addressed in the previous section. Moreover, we can perform a quantitative model comparison using the estimated natural-log evidence. In Table 1, we calculate the natural log-Bayes factors ($\ln(K)$) demonstrating that **S+P** model is decisively preferred (e.g. using the interpretation from Kass & Raftery 1995). It is of note that the $\ln(K)$ value obtained from **Models S+P** and **S** is lower than that obtained from **Models S+P** and **no-glitch**. Since the models are nested, the S+P versus S Bayes factor can be compared to the S versus no-glitch Bayes factor to assess the relative importance of the secular glitch and the step changes in the periodic modulation. Since the former is larger than the latter, this implies that for the spin-down data, the step changes in the modulation period are more significant than the secular changes.

6 COMPARING WITH MODEL-INDEPENDENT VISUALIZATIONS

In Ashton et al. (2017), we introduced a time-period plot to study how the modulation period varies across the observed data span. We now build on this concept in order to understand the implications of the **S+P** model inferences. First, we fit and subtract a first-order polynomial from the raw spin-down rate data. This ensures only the periodic modulations remain, and any information on the average spin-down rate or the second-order spin-down rate is removed. We then plot the Lomb–Scargle periodogram (Lomb 1976; Scargle 1982) applied in a sliding window with a stride length of 1500 d. We varied this stride length, balancing long-duration windows that




Figure 8. Figure similar to Fig. 2 showing the 8(a) spin-down rate data together with the maximum posterior estimate solution of the model; and the 8(b) residuals, for **Model no-glitch**.

reduce the uncertainty on the estimation of the period with short-duration windows that increase the resolution in time.

We find three modes in the spectrum: a primary mode at ~ 460 d (taking a reference epoch of 50000 MJD, the start of the data) and two smaller ones at ~ 250 and ~ 170 d. The 460 and 250 d modes have been modelled in Nijū et al. (2022) and Ashton et al. (2017); the authors of Stairs et al. (2000) stated that there was a strong indication for the presence of a third mode at ~ 167 d and Reznia (2003) confirmed the existence of this third harmonic.

We add to Fig. 10 horizontal lines denoting the epoch of the time parameters obtained by the data fit shown in Fig. 2(a). We also add white dashed lines to represent the three modulation periods and evolution of the three harmonic modes obtained by the fit on the spin-down rate data. They match the modulation periods obtained from the data. Higher derivatives of the modulation period obtained

from the model are not displayed, as the Lomb–Scargle periodogram reveals no additional modes beyond those already presented.

From Fig. 10, we can clearly identify that the modulation period is decreasing over time and continues to do so after the glitch at approximately the same rate. If we compare the modulation period value obtained for the major mode for the first and the last sliding window, we can see that the modulation period has decreased from 489 to 435 d. From the Lomb–Scargle periodogram we can extract the maximum values and perform a linear regression across the entire data set, as well as before and after the glitch. This calculation returned a rate of change of the modulation period of ~ -0.011 s s^{-1} , consistent with the values previously calculated in Ashton et al. (2017) on the shorter pre-glitch data (see Fig. 1). Additionally, we observe that before the glitch, the modulation period changes at a rate of ~ -0.010 s s^{-1} . After the glitch, this rate increases to ~ -0.014 s




Figure 9. Figure similar to Fig. 2 showing the 9(a) spin-down rate data together with the maximum posterior estimate solution of the model; and the 9(b) residuals, for **Model S**.

Table 1. Tabulated \ln Bayes factor, $\ln(K)$, calculated for a comparison between **Model S+P**, and the other models.

Model A	Model B	$\ln(K)$
S+P	no-glitch	1623.60
S+P	S	1486.34

s^{-1} , indicating that the modulation period is decreasing more rapidly post-glitch.

Another noticeable feature is the change of maximum spectral amplitude occurring at the same t_s^η reported by the model, indicated by the line in yellow. A decrease in the spectral amplitude occurs at t_s^η for the first modulation period mode, while the second modulation period mode sees an increase. This is also consistent with the inferences made from **Model S+P**. Finally, a shift in the modulation

period and frequency is observed at t_s^X , matching the model, indicated by the line in red. After this point, the modulation period temporarily shifts to the right, indicating an increase. Subsequently, it resumes its continuous decrease for the remainder of the data range. This behaviour aligns with the model results depicted in Fig. 7.

7 DISCUSSION

In this work, we analyse a newly available high-resolution and expanded data set containing the spin-down rate of PSR B1828–11. The longer data set contains several cycles of observations after the pulsar glitch at 55040.9 MJD. As opposed to previous efforts, in which physics-informed models were developed to explain the behaviour of the pulsar, here we apply a phenomenological model. We considered three models to describe the behaviour of the pulsar: **Model S+P**, which considered the existence of a glitch and changes




Figure 10. Modulation period spectrum of the spin-down rate residuals over a sliding window of 1500 d, on the x-axis, as a function of the mid-point timestamp for each window, on the y-axis. The z-axis shows the Spectral Amplitude. The horizontal lines represent the glitch and step change time parameters (at their exact inferred times) obtained by **Model S+P** and are superimposed upon this figure for comparison purposes. t_s^{ξ} , in blue (solid), represents the glitch time parameter; t_s^{η} , in yellow (dashed), represents the step change time parameter related to a change in amplitude; t_s^{χ} , in red (dotted), represents the step change time parameter related to a change in modulation frequency; and t_s^{δ} , in green (dotted dashed), represents the step change time parameter related to a change in phase offset. The vertical dashed white lines represent three modulation period modes returned by **Model S+P**.

to the periodic modulation of the star; **Model no-glitch**, which assumed that no glitch nor changes to the periodic modulation occurred; and **Model S**, which allowed for a glitch but assumed no modulation changes. **Model S+P** was preferred over the other two, as detailed in Section 5.3, with the Bayes factor values shown in Table 1.

Model S+P was designed to allow for parameter changes at the glitch, with a preference for no change to occur, by using ‘Slab-and-spike’ priors. For each component of the model, i.e. the spin-down rate, the amplitude, the phase-offset and the modulation frequency, we allowed for a separate parameter recording the time at which the step occurred, to see if these changes coincided with the glitch time. This was not the case, with step changes for some parameters being found to occur before the glitch: a decrease of modulation amplitude, described by $\eta_1 = -0.175 \pm 0.008$ was observed for the main modulation period mode at t_s^{η} ; a change in modulation frequency was observed, as seen in Fig. 7, at t_s^{χ} ; all cosine components observed a change in the phase-offset at t_s^{δ} . Additionally, we found that the spin-down rate experienced an exponential recovery after the glitch t_s^{ξ} .

To add validity to this model, the data was analysed without a model, as described in Section 6. We obtained a Lomb–Scargle periodogram spectrum from the data. This shows that the modulation period of the pulsar is decreasing at a rate of $\sim -0.010 \text{ s s}^{-1}$ before

the glitch and decreasing more rapidly after the glitch ($\sim -0.014 \text{ s s}^{-1}$); the modulation period modes experienced spectral amplitude changes before the glitch time; and we visually identify a step change to the modulation frequency and thus the modulation period, also before the glitch.

7.1 Interpreting the results in light of astrophysical models

7.1.1 Planetary companions

It is clear already from the correlated changes in the spin-down and pulse shape that a planetary explanation for the modulations is unlikely. Nevertheless, the observation of a decreasing modulation period (Ashton et al. 2017) invites reconsideration of the planetary explanation, with some coupling torque between the star and planet(s) yielding the inspiral and explaining the changing modulation period. However, while no generative model exists, we feel that any such model would necessarily require the modulations to arise from the smoothly varying orbital separation of the planets and star. Therefore, the observation in this work is that there are distinct instantaneous changes in the modulation seems difficult to explain with a planetary hypothesis. Moreover, the observation that there are up to eight harmonically related sinusoids would also call for up to

eight planets, following the original arguments that the two sinusoids correspond to two planets. Once again, this feels implausible.

7.1.2 Free precession

The global difficulties of reconciling free precession, the decreasing modulation period, and the glitch have already been discussed in Jones et al. (2017). Here, we can quantify the observation from Stairs et al. (2019) and Shaw et al. (2022) that contrary to some of the predictions of Jones et al. (2017), the modulations continue after the glitch, constraining the models connecting the interior of the star to the cause of the modulations. Furthermore, the observation that there are distinct changes (at different times) in the modulation amplitude, frequency, and phase adds to the challenge of interpreting this event in light of free precession. However, the observation of multiple harmonically related sinusoids offers a new opportunity to test the model. Namely, in Jones & Andersson (2001), the precession model used in later works was developed with an expansion in the small angle θ between the symmetry axis of the (assumed biaxial) moment of inertia tensor and the angular momentum. Therefore, a more physically accurate model can be obtained by either expanding the model to include higher-order terms. Since the precession model has no additional degrees of freedom left, it will be interesting to discover if the amplitude coefficients of the harmonically related sinusoids measured here are consistent with the predictions of precession, enabling a new test of precession. However, to explain the observed beamwidth data, such a model would need to be extended as suggested by Stairs et al. (2019) to combine the long-term precession behaviour with quantized profile switches.

7.1.3 Magnetospheric switching

There is no single well-defined magnetospheric switching model: in a sense, it is a set of observations rather than a generative model itself. It is, therefore, not straightforward to connect our observations to such a model. Moreover, since we are using only the spin-down rate data, we are insensitive to the rapid switches between states and can only discuss the long-time-scale periodic modulations in this switching rate. Nevertheless, our phenomenological study reveals several insights into whatever process drives this. First, there is decisive evidence for more than two terms in the harmonic expansion; this is observable directly by our fit to the data but also by the non-zero posterior support for these terms. The amplitude of these terms could provide a way to test mechanisms for the clock (e.g. as proposed for precession in the previous section). Second, we find evidence for distinct changes in the spectral amplitude of the harmonic sinusoid, with a sudden shift from the fundamental to the first harmonic. This is interesting as it suggests another variability mechanism for the periodic modulations. It would be interesting to study the raw data of PSR B1828–11 to identify if there are corresponding systematic changes in the beam shape during this transition (or, indeed, any of the observed step changes).

In Seymour & Lorimer (2013), the authors introduced evidence that PSR B1828–11 was displaying chaotic behaviour consistent with a system with three governing variables. From this work and further discussion (Stairs et al. 2019), it is proposed that the spin-down rate and mode transition rate act as two of the governing variables. However, it is unclear to see how this observation can be connected to a physical process to predict the observed chaotic behaviour.

8 OUTLOOK

The high-quality data released by Niřu et al. (2022) has enabled a new and detailed study of PSR B1828–11 using the inferred spin-down rate. Since the end of the data set studied in this work, JBO has continued observing PSR B1828–11 and we expect there to be several more cycles to study. Moreover, there are other pulsars which display similar (if less clear) behaviours. For example the authors of Zubieta et al. (2024) have reported on changes to the amplitude and frequency of PSR J0742–2822 following a glitch. We believe the tools and techniques developed in this work could be applied to larger data sets, with the ultimate goal of providing quantitative measures of the behaviour to help us constrain models. However, one key missing aspect is that we are studying only the spin-down rate and neglecting information about the mode-transition rate. Therefore, we believe further methodological work is needed to develop approaches that can automate the analysis of pulsars. This would allow the study of both their rapidly changing beam shape and their long-term timing properties. We believe this has the capacity to answer long-held questions about the star’s interaction with its magnetosphere.

The work presented here shows that a model that allows for sudden step changes in the amplitude, phase, and frequency parameters in the long-term periodicity in the spin-down of PSR B1828–11 is a better fit to the data than a model that does not allow for sudden step changes in these parameters. Significantly, the model locates these step changes at *three different times*, all well *before* the glitch itself. This is somewhat surprising, and difficult to account for in terms of a physical model. This suggests that other models of the long-term periodicity, not based on sudden step changes, may also be worth exploring.

ACKNOWLEDGEMENTS

We utilize the NUMPY (Harris et al. 2020) and MATPLOTLIB (Hunter 2007) libraries for data processing and visualization and the SCIPY library (Virtanen et al. 2020) for implementation of the Lomb–Scargle periodogram. DIJ acknowledges support from the UK Research and Innovation (UKRI) Science and Technology Funding Council (STFC) via grant No. ST/R00045X/1.

DATA AVAILABILITY

The data used in this publication are available in Keith & Niřu (2023a). The code needed to produce the results shown in this publication can be shared upon request to the corresponding author.

REFERENCES

- Akgün T., Link B., Wasserman I., 2006, *MNRAS*, 365, 653
- Ashton G., Jones D. I., Prix R., 2016, *MNRAS*, 458, 881
- Ashton G., Jones D. I., Prix R., 2017, *MNRAS*, 467, 164
- Ashton G., et al., 2019, *ApJS*, 241, 27
- Bailes M., Lyne A., Shemar S., 1993, in Phillips J., Thorsett S., Kulkarni S., eds, ASP Conf. Ser. Vol. 36, Planets Around Pulsars. Astron. Soc. Pac., San Francisco, p. 19
- Basu A. et al., 2021, *MNRAS*, 510, 4049
- Brook P. R., Karastergiou A., Johnston S., Kerr M., Shannon R. M., Roberts S. J., 2016, *MNRAS*, 456, 1374
- Carlin J. B., Dunson D. B., Gelman A., Rubin D. B., Stern H. S., Vehtari A., 2013, Bayesian Data Analysis, 3rd edn. CRC Press, New York
- Green P. J., 2003, Oxford Statistical Science Series, 1, 179
- Harris C. R. et al., 2020, *Nature*, 585, 357

Haskell B., Jones D. I., 2024, *Astrophys. J.*, 157, 102921
Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655
Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
Jones D. I., 2012, *MNRAS*, 420, 2325
Jones D. I., Andersson N., 2001, *MNRAS*, 324, 811
Jones D. I., Ashton G., Prix R., 2017, *Phys. Rev. Lett.*, 118, 261101
Kass R. E., Raftery A. E., 1995, *J. Am. Stat. Assoc.*, 90, 773
Keith M. J., Nițu I. C., 2023a, Quasi-periodic spin down variations in 17 pulsars, Zenodo, doi:10.5281/zenodo.7664166
Keith M. J., Nițu I. C., 2023b, *MNRAS*, 523, 4603
Koposov S. et al., 2024, joshspiegel/dynesty: v2.1.4, Zenodo, doi:10.5281/zenodo.12537467
Kramer M., Lyne A. G., O’Brien J. T., Jordan C. A., Lorimer D. R., 2006, *Science*, 312, 549
Link B., Epstein R. I., 2001, *ApJ*, 556, 392
Liu K., Yue Y. L., Xu R. X., 2007, *MNRAS*, 381, L1
Lomb N. R., 1976, *Ap&SS*, 39, 447
Lorimer D., Kramer M., 2005, *Handbook of Pulsar Astronomy*. Cambridge Observing Handbooks for Research Astronomers, Cambridge Univ. Press, Cambridge
Lovell A. C. B., 1957, *Nature*, 180, 60
Lower M. E. et al. 2025, *MNRAS*, preprint (arXiv:2501.03500)
Lyne A., Hobbs G., Kramer M., Stairs I., Stappers B., 2010, *Science*, 329, 408
Malsiner-Walli G., Wagner H., 2016, *Austrian J. Stat.*, 40, 241
Nițu I. C., Keith M. J., Stappers B. W., Lyne A. G., Mickaliger M. B., 2022, *MNRAS*, 512, 2446
Parthasarathy A. et al., 2019, *MNRAS*, 489, 3810
Perera B. B. P., Stappers B. W., Weltevredre P., Lyne A. G., Bassa C. G., 2015, *MNRAS*, 446, 1380
Rezanian V., 2003, *A&A*, 399, 653
Scargle J. D., 1982, *ApJ*, 263, 835
Seymour A. D., Lorimer D. R., 2013, *MNRAS*, 428, 983
Shaham J., 1977, *ApJ*, 214, 251
Shaw B. et al., 2022, *MNRAS*, 513, 5861
Skilling J., 2004, in Fischer R., Preuss R., von Toussaint U., eds, AIP Conf. Proc. Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering. Am. Inst. Phys., New York, p. 395
Speagle J. S., 2020, *MNRAS*, 493, 3132
Stairs I. H., Lyne A. G., Shemar S. L., 2000, *Nature*, 406, 484
Stairs I. H., Athanasiadis D., Kramer M., Lyne A. G., 2003, in Bailes M., Nice D. J., Thorsett S. E., eds, ASP Conf. Ser. Vol 302, Radio Pulsars. Astron. Soc. Pac., San Francisco, p. 249
Stairs I. H. et al., 2019, *MNRAS*, 485, 3230
Tong H. et al., 2025, *ApJS*, 276, 50
Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
Zubieta E. et al. 2024, *A&A*, 694, A124, preprint (arXiv:2412.17766)

APPENDIX A: PRIOR DISTRIBUTIONS AND SUMMARY STATISTICS

Table A1. Prior distributions for the **Model S+P**’s parameters. Parameters with priors denoted with ‘SS’ have slab-spike priors applied to them.

	Prior	Units
$\dot{\nu}_0$	$\text{Unif}(-2.74 \times 10^{-3}, -2.72 \times 10^{-3})$	d^{-2}
$\dot{\nu}_1$	$\text{Unif}(-2.73 \times 10^{-7}, 2.73 \times 10^{-7})$	d^{-3}
$\dot{\nu}_2$	$\text{Unif}(-2.73 \times 10^{-11}, 2.73 \times 10^{-11})$	d^{-4}
$\dot{\nu}_3$	$\text{Unif}(-2.73 \times 10^{-15}, 2.73 \times 10^{-15})$	d^{-5}
τ	$\text{Unif}(0, 500)$	d
t_s^{ξ}	$\text{Unif}(54990.90, 55090.90)$	d
ξ_p^0	$\text{SS} + \text{Unif}(-0.01, 0.01)$	d
ξ_l^0	$\text{SS} + \text{Unif}(-0.01, 0.01)$	d
A_{1-9}	$\text{SS} + \text{Unif}(0, 1.00 \times 10^{-5})$	–
η_{1-9}	$\text{SS} + \mathcal{N}(0, 1)$	–
$\Delta\phi_{1-9}$	$\text{Unif}(-\pi, \pi)$	rad
δ_{1-9}	$\text{SS} + \mathcal{N}(0, 1)$	–
f_0	$\text{SS} + \text{Unif}(2.11 \times 10^{-3}, 2.33 \times 10^{-3})$	d^{-1}
f_1	$\text{SS} + \text{Unif}(-2.22 \times 10^{-7}, 2.22 \times 10^{-7})$	d^{-2}
f_2	$\text{SS} + \text{Unif}(-2.22 \times 10^{-11}, 2.22 \times 10^{-11})$	d^{-3}
χ_0	$\text{SS} + \text{Unif}(-1, 1)$	–
χ_1	$\text{SS} + \text{Unif}(-1, 1)$	–
χ_2	$\text{SS} + \text{Unif}(-1, 1)$	–
Δt_s^{χ}	$\text{SS} + \text{Unif}(-5000, 2000)$	–
Δt_s^{δ}	$\text{SS} + \text{Unif}(-5000, 2000)$	–
Δt_s^{η}	$\text{SS} + \text{Unif}(-5000, 2000)$	–

Table A2. Prior distributions for **Model no-glitch**’s parameters. Parameters with priors denoted with ‘SS’ have slab-spike priors applied to them.

	Prior	Units
$\dot{\nu}_0$	$\text{Unif}(-2.74 \times 10^{-3}, -2.72 \times 10^{-3})$	d^{-2}
$\dot{\nu}_1$	$\text{Unif}(-2.73 \times 10^{-7}, 2.73 \times 10^{-7})$	d^{-3}
$\dot{\nu}_2$	$\text{Unif}(-2.73 \times 10^{-11}, 2.73 \times 10^{-11})$	d^{-4}
$\dot{\nu}_3$	$\text{Unif}(-2.73 \times 10^{-15}, 2.73 \times 10^{-15})$	d^{-5}
A_{1-9}	$\text{SS} + \text{Unif}(0, 1.00 \times 10^{-5})$	–
$\Delta\phi_{1-9}$	$\text{Unif}(-\pi, \pi)$	rad
f_0	$\text{SS} + \text{Unif}(2.11 \times 10^{-3}, 2.33 \times 10^{-3})$	d^{-1}
f_1	$\text{SS} + \text{Unif}(-2.22 \times 10^{-7}, 2.22 \times 10^{-7})$	d^{-2}
f_2	$\text{SS} + \text{Unif}(-2.22 \times 10^{-11}, 2.22 \times 10^{-11})$	d^{-3}

Table A3. Prior distributions for **Model S**’s parameters. Parameters with priors denoted with ‘SS + Unif’ have slab-spike priors applied to them.

	Prior	Units
$\dot{\nu}_0$	$\text{Unif}(-2.74 \times 10^{-3}, -2.72 \times 10^{-3})$	d^{-2}
$\dot{\nu}_1$	$\text{Unif}(-2.73 \times 10^{-7}, 2.73 \times 10^{-7})$	d^{-3}
$\dot{\nu}_2$	$\text{Unif}(-2.73 \times 10^{-11}, 2.73 \times 10^{-11})$	d^{-4}
$\dot{\nu}_3$	$\text{Unif}(-2.73 \times 10^{-15}, 2.73 \times 10^{-15})$	d^{-5}
τ	$\text{Unif}(0, 500)$	d
t_s^{ξ}	$\text{Unif}(54990.90, 55090.90)$	d
ξ_p^0	$\text{SS} + \text{Unif}(-0.01, 0.01)$	d
ξ_l^0	$\text{SS} + \text{Unif}(-0.01, 0.01)$	d
A_{1-9}	$\text{SS} + \text{Unif}(0, 1.00 \times 10^{-5})$	–
$\Delta\phi_{1-9}$	$\text{Unif}(-\pi, \pi)$	rad
f_0	$\text{SS} + \text{Unif}(2.11 \times 10^{-3}, 2.33 \times 10^{-3})$	d^{-1}
f_1	$\text{SS} + \text{Unif}(-2.22 \times 10^{-7}, 2.22 \times 10^{-7})$	d^{-2}
f_2	$\text{SS} + \text{Unif}(-2.22 \times 10^{-11}, 2.22 \times 10^{-11})$	d^{-3}

Table A4. Maximum posterior distribution summary, with their standard deviations, for **Model S + P**'s parameters.

	Posterior median(s.d.)	Units
\dot{v}_0	$-2.72739(5) \times 10^{-3}$	d^{-2}
\dot{v}_1	$9.0(3) \times 10^{-10}$	d^{-3}
\dot{v}_2	$6(2) \times 10^{-14}$	d^{-4}
\dot{v}_3	$-2(1) \times 10^{-17}$	d^{-5}
τ	121(14)	d
t_s^ξ	55047(3)	d
ξ_p^0	$-5(3) \times 10^{-5}$	d
ξ_t^0	$1.23(8) \times 10^{-3}$	d
A_1	$5.77(3) \times 10^{-6}$	—
A_2	$4.77(3) \times 10^{-6}$	—
A_3	$2.03(3) \times 10^{-6}$	—
A_4	$9.3(3) \times 10^{-7}$	—
A_5	$7.4(3) \times 10^{-7}$	—
A_6	$2.0(3) \times 10^{-7}$	—
A_7	$5.8(3) \times 10^{-7}$	—
A_8	$2.6(3) \times 10^{-7}$	—
A_9	$0(1) \times 10^{-9}$	—
η_1	$-0.175(8)$	—
η_2	0.19(1)	—
η_3	$-0.38(2)$	—
η_4	$-0.32(5)$	—
η_5	0(3)	—
η_6	0.8(4)	—
η_7	$-0.30(9)$	—
η_8	$-0.8(1)$	—
η_9	$-0.03(7)$	—
$\Delta\phi_1$	2.30(1)	rad
$\Delta\phi_2$	$-1.48(2)$	rad
$\Delta\phi_3$	3.139(9)	rad
$\Delta\phi_4$	$-0.32(3)$	rad
$\Delta\phi_5$	1.4(1)	rad
$\Delta\phi_6$	$-0.6(2)$	rad
$\Delta\phi_7$	3.13(4)	rad
$\Delta\phi_8$	2.7(1)	rad
$\Delta\phi_9$	2(2)	rad
δ_1	$-0.143(5)$	—
δ_2	0.39(1)	—
δ_3	$-0.276(6)$	—
δ_4	4.4(4)	—
δ_5	$-0.59(5)$	—
δ_6	3.7(5)	—
δ_7	$-0.47(2)$	—
δ_8	0(4)	—
δ_9	1.5(7)	—
f_0	$2.1748(3) \times 10^{-3}$	d^{-1}
f_1	$1.847(7) \times 10^{-8}$	d^{-2}
f_2	$0(6) \times 10^{-16}$	d^{-3}
χ_0	$3.9(2) \times 10^{-4}$	—
χ_1	0.81(2)	—
χ_2	0.2(4)	—
Δt_s^χ	$-1424(4)$	—
Δt_s^δ	$-4420(20)$	—
Δt_s^η	$-740(6)$	—

Table A5. Maximum posterior distribution summary, with their standard deviations, for **Model no-glitch**'s parameters.

	Posterior median(s.d.)	Units
\dot{v}_0	$-2.72754(4) \times 10^{-3}$	d^{-2}
\dot{v}_1	$9.5(2) \times 10^{-10}$	d^{-3}
\dot{v}_2	$1.8(2) \times 10^{-13}$	d^{-4}
\dot{v}_3	$3(1) \times 10^{-17}$	d^{-5}
A_1	$5.27(3) \times 10^{-6}$	—
A_2	$5.08(3) \times 10^{-6}$	—
A_3	$1.66(3) \times 10^{-6}$	—
A_4	$7.1(3) \times 10^{-7}$	—
A_5	$6.6(3) \times 10^{-7}$	—
A_6	$2.5(3) \times 10^{-7}$	—
A_7	$3.7(3) \times 10^{-7}$	—
A_8	$9(5) \times 10^{-8}$	—
A_9	$0(3) \times 10^{-9}$	—
$\Delta\phi_1$	2.070(7)	rad
$\Delta\phi_2$	$-1.875(8)$	rad
$\Delta\phi_3$	2.53(2)	rad
$\Delta\phi_4$	$-1.16(5)$	rad
$\Delta\phi_5$	0.85(5)	rad
$\Delta\phi_6$	$-1.6(1)$	rad
$\Delta\phi_7$	2.6(1)	rad
$\Delta\phi_8$	3(2)	rad
$\Delta\phi_9$	2(2)	rad
f_0	$2.1874(3) \times 10^{-3}$	d^{-1}
f_1	$2.16(1) \times 10^{-8}$	d^{-2}
f_2	$0(1) \times 10^{-15}$	d^{-3}

Table A6. Maximum posterior distribution summary, with their standard deviations, for **Model S**'s parameters.

	Posterior median(s.d.)	Units
\dot{v}_0	$-2.72737(7) \times 10^{-3}$	d^{-2}
\dot{v}_1	$9.2(4) \times 10^{-10}$	d^{-3}
\dot{v}_2	$9(2) \times 10^{-14}$	d^{-4}
\dot{v}_3	$-0.03(1) \times 10^{-17}$	d^{-5}
τ	74(8)	d
t_s^ξ	55091(2)	d
ξ_p^0	$-3.67(4) \times 10^{-5}$	d
ξ_t^0	$2.1037(1) \times 10^{-3}$	d
A_1	$5.32(3) \times 10^{-6}$	—
A_2	$5.07(3) \times 10^{-6}$	—
A_3	$1.67(3) \times 10^{-6}$	—
A_4	$6.8(3) \times 10^{-7}$	—
A_5	$6.7(3) \times 10^{-7}$	—
A_6	$2.3(3) \times 10^{-7}$	—
A_7	$3.6(3) \times 10^{-7}$	—
A_8	$1.4(6) \times 10^{-7}$	—
A_9	$0(4) \times 10^{-9}$	—
$\Delta\phi_1$	2.088(6)	rad
$\Delta\phi_2$	$-1.853(7)$	rad
$\Delta\phi_3$	2.57(2)	rad
$\Delta\phi_4$	$-1.11(5)$	rad
$\Delta\phi_5$	0.88(5)	rad
$\Delta\phi_6$	$-1.7(1)$	rad
$\Delta\phi_7$	2.52(9)	rad
$\Delta\phi_8$	3(2)	rad
$\Delta\phi_9$	2(2)	rad
f_0	$2.1868(3) \times 10^{-3}$	d^{-1}
f_1	$2.14(1) \times 10^{-8}$	d^{-2}
f_2	$0(1) \times 10^{-16}$	d^{-3}

This paper has been typeset from a \LaTeX file prepared by the author.