



Centre for
Astrophysics and
Supercomputing

Exploring the magnetospheric and rotational properties of radio pulsars

Marcus Edward Lower

Presented in fulfillment of the requirements
of the degree of Doctor of Philosophy

January 2022

School of Science, Computing and Engineering Technologies
Swinburne University of Technology

“Magnetars can do anything”
– John Reynolds

Abstract

Radio pulsars are rotating neutron stars that emit beams of coherent electromagnetic radiation often aligned with their magnetic poles. When these beams cross our line of sight, we can detect them as highly periodic pulsed radio sources. The propagation of their radio pulses through intervening material can introduce a variety of frequency and polarisation dependent phenomena. Modelling of these propagation effects allow us to probe not only the properties of the ionised interstellar medium between us and the pulsar, but also the magneto-ionic environments that exists within the immediate vicinity of neutron stars. Additionally, the regularity at which their radio pulses arrive at the Earth allows us to perform tests of fundamental physics and undertake studies of nuclear matter in extreme conditions through precision pulsar timing. In this thesis we demonstrate how increases in observing cadence, bandwidth and sensitivity afforded by the latest generation of radio telescopes can be used to explore the rotational and magnetospheric properties of these objects.

We used modern pulsar timing methods and inference techniques to assess the prevalence of long-term rotational irregularities across a large sample of 300 pulsars that are observed every few days with the Molonglo Observatory Synthesis Telescope. Although we employed a relatively simple model to infer the scaling of timing noise strength with pulsar spin-frequencies and spin-down rates, the hierarchical inference framework that we developed can be easily adapted to accommodate more complex, astrophysical population models. These methods were also applied to a set of 74 pulsars that have been observed by the Parkes radio telescope over 5-30 years. We detected 124 transient spin-up events in these pulsars (glitches), of which 74 were new discoveries. A newly developed hidden-Markov based glitch detection algorithm verified the resulting glitch catalogue was sample complete. This subsequently allowed us to explore the inter-glitch rotational evolution of 33 pulsars that had measurable braking indices.

We also investigated the spectropolarimetric properties of Swift J1818.0–1607, a radio-loud magnetar that was discovered in March 2020. Observations taken soon after its initial discovery using the Parkes Ultra-Wideband Low (UWL) receiver system revealed the pulsed radio emission from the magnetar more closely resembled those detected from more ordinary, rotation-powered pulsars, as opposed to other radio-loud magnetars. This led us to speculate that Swift J1818.0–1607 may be represent a rare class of transitional objects that followed a non-standard evolutionary path. Our continued observations of Swift J1818.0–1607 with the UWL revealed it possesses a highly dynamic magnetosphere.

A broad range of radio emission phenomena seldom seen among other pulsars were detected in the radio pulses emitted by the magnetar. Fitting of a rotating vector model to the linear polarisation position angle suggest it is an orthogonal rotator with two co-located magnetic poles positioned near its equator.

Finally, we analysed a set of high sensitivity observations taken by MeerKAT of the magnetospheric eclipses detected in the Double Pulsar system PSR J0737–3039A/B. Our modelling of variations in the eclipse light curves recovered a geodetic precession rate of pulsar B is consistent with Einstein’s theory of General Relativity at the 7.4% level. This represents an almost two-fold improvement over previous measurements with this system. We also detected significant conversion of linear to circular polarisation during eclipse ingress, superior conjunction and egress. This may be the first direct evidence the plasma found within the closed-field region of neutron star magnetospheres are comprised of highly relativistic electron-positron pairs.

Acknowledgements

While a PhD is a qualification that can only go to an individual, it is by no means an individual effort to complete one. I am incredibly grateful to have worked alongside and been supported by many wonderful people during the past three-and-a-bit years. First of all, I would like to thank my PhD supervisors, Matthew, Ryan and Simon. I wouldn't have become the scientist I am today without your boundless enthusiasm for all things pulsars, occasional pearls of wisdom and general support over the past three years. I am also lucky to have had the privilege of collaborating with many other scientists on various projects throughout this time. In particular, I thank Liam, Fernando, Michael, Shi, Hao, Chris, Andrew, Dick, René and Norbert, for their contributions to both work that went into this thesis and elsewhere. And of course, I must also give thanks to the many friends and colleagues in OzGrav, the Swinburne Pulsar/FRB/GW group, the wider Swinburne PhD student cohort and the Monash GW group. The past three years wouldn't have been nearly as enjoyable without working alongside and hanging out with the fantastic people in each of these groups.

I thoroughly enjoyed my visits to the CSIRO Marsfield site, thanks largely to the wonderful hospitality of the ATNF staff. I am also immensely grateful for the hard work put in by Molonglo team (particularly Chris Flynn and Tim Bateman), the Parkes Observatory support staff (special shout out to John Sarkissian) and the telescope operators at SRAO (especially Sarah Buchner), both in maintaining the telescopes and helping out when things went horribly wrong.

Finally, thank you to all the wonderful and supportive people in my life who encouraged me throughout my studies. In particular, my parents, Astrid and Richard, and my brother Michael. A special thanks goes to Laura. For putting up with my many rants about pulsars and magnetars, early-morning observing, reading paper drafts and forcing me to step away from my work when needed.

The research presented throughout this thesis was supported by Australian Government Research Training Program, the Australian Research Council Laureate Fellowship FL150100148, and CSIRO Space and Astronomy.

Declaration

The work presented in this thesis has been carried out in the Centre for Astrophysics & Supercomputing at Swinburne University of Technology and the Australia Telescope National Facility at CSIRO Space and Astronomy between 2018 and 2021. This thesis contains no material that has been accepted for the award of any other degree or diploma. To the best of my knowledge, this thesis contains no material previously published or written by another author, except where due reference is made in the text of the thesis. The content of the chapters listed below has either appeared in or been submitted to refereed journals. Minor alterations have been made to the published papers in order to maintain argument continuity and consistency of spelling and style.

- Chapter 3 was published in *Monthly Notices of the Royal Astronomical Society* as “The UTMOST pulsar timing programme – II. Timing noise across the pulsar population”, authored by **M. E. Lower**, M. Bailes, R. M. Shannon, S. Johnston, C. Flynn, S. Osłowski, V. Gupta, W. Farah, T. Bateman, A. J. Green, R. Hunstead, A. Jameson, F. Jankowski, A. Parthasarathy, D. C. Price, A. Sutherland, D. Temby and V. Venkatraman Krishnan.

I led the project and did most of the software development and data analysis directly required for the manuscript. M. Bailes, R. M. Shannon, S. Johnston, C. Flynn, S. Osłowski, V. Gupta, W. Farah, R. Hunstead, F. Jankowski, D. C. Price and V. Venkatraman Krishnan each provided comments on the manuscript prior to submission. T. Bateman is the site Senior Technical Officer who oversaw the the daily operations at UTMOST. The remaining co-authors performed miscellaneous tasks on the UTMOST project. The comments from the anonymous reviewer during the publication phase were also helpful in improving the overall quality and presentation of the manuscript.

- Chapter 4 was published in *Monthly Notices of the Royal Astronomical Society* as “The impact of glitches on young pulsar rotational evolution”, authored by **M. E. Lower**, S. Johnston, L. Dunn, R. M. Shannon, M. Bailes, S. Dai, M. Kerr, R. N. Manchester, A. Melatos, L. S. Oswald, A. Parthasarathy, C. Sobey and P. Weltevrede.

I led the project and performed most of the software development and data analysis directly required for the manuscript. L. Dunn ran the glitch detection algorithm on the data and derived the glitch upper-limits. S. Johnston was instrumental in provid-

ing valuable suggestions throughout the duration of the project. I wrote most of the manuscript with help from S. Johnston, L. Dunn R. M. Shannon, with suggestions from the remaining co-authors. The comments from the anonymous referee assisted in improving the overall quality of the final manuscript. Most of the co-authors helped perform the observations conducted with the 64-m Parkes radio telescope.

- Chapter 5 was published in *The Astrophysical Journal Letters* as “Spectropolarimetric properties of Swift J1818.0–1607: A 1.4s radio magnetar”, authored by **M. E. Lower**, R. M. Shannon, S. Johnston and M. Bailes.

I led the data collection, processing, software development and analysis required for this manuscript. M. Bailes provided software for computing pulse widths and signal-to-noise ratios. R. M. Shannon, S. Johnston and M. Bailes each contributed comments and suggestions on the manuscript. The comments and corrections from the anonymous referee assisted in improving the overall quality of the final manuscript during the publication process.

- Chapter 6 was published in *Monthly Notices of the Royal Astronomical Society* as “The dynamic magnetosphere of Swift J1818.0–1607”, authored by **M. E. Lower**, S. Johnston, R. M. Shannon, M. Bailes and F. Camilo.

I led the data collection, processing, software development and analysis required for this manuscript. S. Johnston and R. M. Shannon assisted in interpreting the data and double-checking results. S. Johnston, R. M. Shannon and M. Bailes each provided comments on the manuscript. F. Camilo is the principle investigator of the Parkes magnetar monitoring programme. The anonymous referee provided minor comments that improved the final published manuscript.

The candidate also made contributions to the following publications during their candidature:

- P. Kumar, R. M. Shannon, **M. E. Lower**, et al., 2021, “Circularly polarized radio emission from the repeating fast radio burst source FRB 20201124A” submitted to *Monthly Notices of the Royal Astronomical Society* arXiv:2109.11535 – 3rd author, contributed software for data analysis and interpretation.
- S. Johnston and **M. E. Lower**, 2021, “A supernova remnant association for the fast-moving pulsar PSR J0908–4913” *Monthly Notices of the Royal Astronomical Society Letters* 507 L41 – 2nd author, contributed to data analysis and interpretation, co-wrote the manuscript with S. Johnston.

- B. Goncharov, R. M. Shannon, D. J. Reardon, et al., 2021, “On the evidence for a common-spectrum process in the search for the nanohertz gravitational wave background with the Parkes Pulsar Timing Array” *The Astrophysical Journal Letters* 917 L19 – 10th author, contributed as a member of the PPTA.
- L. Dunn, **M. E. Lower** and A. Melatos, 2021, “Effects of periodicity in observation scheduling on parameter estimation of pulsar glitches” *Monthly Notices of the Royal Astronomical Society* 504 3399 – 2nd author, contributed data analysis, interpretation, and extensive comments on the manuscript.
- M. Kramer, I. Stairs, V. Venkatraman-Krishnan, et al., 2021, “The Relativistic Binary Programme on MeerKAT: Science objectives and first results” *Monthly Notices of the Royal Astronomical Society* 504 2094 – 16th author, contributed data analysis, Parkes observations as a MeerTime Rel. Bin. group member.
- M. Bailes, C. G. Bassa, G. Bernardi, et al., 2021 “Multi-frequency observations of SGR J1935+2154” *Monthly Notices of the Royal Astronomical Society* 503 5367 – 37th author, contributed to data analysis.
- I. Romero-Shaw, C. Talbot, S. Biscoveanu, et al., 2020 “Bayesian inference for compact binary coalescences with Bilby: Validation and application to the first LIGO-Virgo gravitational-wave transient catalogue” *Monthly Notices of the Royal Astronomical Society* 499 3295 – 41st author, contributed as a core Bilby developer.
- H. Ding, A. T. Deller, **M. E. Lower**, et al., 2020, “A magnetar parallax” *Monthly Notices of the Royal Astronomical Society* 498 3736 – 3rd author, contributed observations at Parkes, data analysis, interpretation and manuscript comments.
- M. Bailes, A. Jameson, F. Abbate, et al., 2020, “The MeerKAT Telescope as a Pulsar Facility: System verification and early science results from MeerTime” *Publications of the Astronomical Society of Australia* 37 e028 – 27th author, contributed as a member of MeerTime.
- M. Kerr, D. J. Reardon, G. Hobbs, et al. “The Parkes Pulsar Timing Array Project: Second data release” 37 e020 – 23rd author, contributed as a member of the PPTA.
- “The UTMOST Survey for Magnetars, Intermittent pulsars, RRATs and FRBs I: System description and overview” *Monthly Notices of the Royal Astronomical Society* 492 4752 – 19th author, contributed as an UTMOST project member.

-
- W. Farah, C. Flynn, M. Bailes, et al., 2019, “Five new real-time detections of Fast Radio Bursts with UTMOST” *Monthly Notices of the Royal Astronomical Society*, 488 2989 – 12th author, contributed as an UTMOST project member.
 - E. V. Gotthelf, J. P. Halpern, J. A. J. Alford, et al., 2019 “The 2018 X-ray and Radio Outburst of Magnetar XTE J1810–197” *The Astrophysical Journal Letters*, 874 L25 – 8th author, contributions to data analysis and interpretation.
 - S. Dai, **M. E. Lower**, M. Bailes, et al., 2019, “Wideband polarized radio emission from the newly revived radio magnetar XTE J1810–197”, *The Astrophysical Journal Letters*, 874 L14 – 2nd author, data analyst and co-wrote the manuscript with S. Dai.
 - G. Ashton, M. Hübner, P. D. Lasky, et al., 2019, “Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy”, *The Astrophysical Journal Supplementary Series*, 241 27 – 13th author, contributed as a Bilby core developer.
 - **M. E. Lower**, E. Thrane, P. D. Lasky et al., 2019, “Measuring eccentricity in binary black hole inspirals with gravitational waves” *Physical Review D* 98, 083028 – 1st author, I led the study and wrote the manuscript.

These articles do not form a part of this thesis.



Marcus Edward Lower
Melbourne, Victoria, Australia
January 2022

Contents

Abstract	ii
Acknowledgements	v
Declaration	vii
List of Figures	xvi
List of Tables	xviii
1 Introduction	1
1.1 Neutron stars and pulsars	1
1.1.1 Internal structure	3
1.1.2 General properties	4
1.1.3 Polarimetry	5
1.2 Propagation effects	8
1.2.1 Dispersion	8
1.2.2 Scattering and scintillation	9
1.2.3 Faraday rotation	10
1.3 The neutron star zoo	13
1.3.1 Pulsar birth, life and death	13
1.3.2 Pulsar re-birth via recycling	16
1.3.3 Magnetars	18
1.4 Rotational irregularities in pulsars	20
1.4.1 White noise and profile jitter	21
1.4.2 Red timing noise	21
1.4.3 Glitches	22
1.5 The Double Pulsar PSR J0737–3039A/B	23
1.5.1 Magnetospheric eclipses	23
1.6 Thesis outline	24
2 Observing pulsars with radio telescopes	27
2.1 Basic signal path	27
2.2 The Parkes radio telescope: A single-dish instrument	31
2.3 UTMOST and MeerKAT: Interferometers	33
2.4 Pulsar timing	36

2.4.1	Binary pulsars	38
3	Timing noise across the pulsar population	43
3.1	Introduction	43
3.2	Observations	46
3.2.1	System overview	46
3.2.2	Radio frequency interference	46
3.2.3	Pulsar-timing dataset	47
3.3	Pulsar-timing strategy	48
3.3.1	Phenomenological timing noise model	50
3.3.2	Bayesian framework	50
3.3.3	Braking indices	53
3.3.4	Glitch parameter estimation	54
3.4	Results	55
3.4.1	Updated timing models	55
3.4.2	Spin frequency second derivatives and braking indices	55
3.4.3	Red timing noise properties	59
3.4.4	Pulsar glitch reanalysis	61
3.5	Discussion	68
3.5.1	Quantifying timing noise strength	68
3.5.2	Correlations with individual pulsar properties	70
3.5.3	Scaling relation fitting and hyper-parameter estimation	72
3.5.4	Two noteworthy pulsars	77
3.6	Conclusions	79
4	The impact of glitches on the rotational evolution of young pulsars	83
4.1	Introduction	83
4.2	Observations	86
4.3	Pulsar inference framework	87
4.3.1	Glitch detection and phase connection	87
4.3.2	Measuring glitch properties	88
4.3.3	Pulsar properties and extended noise models	90
4.4	Glitches and timing results	92
4.4.1	Minimum detectable glitch size and sample completeness	95
4.4.2	Vela pulsar timing and pulsars with more than five glitches	97
4.4.3	Recoveries	99

4.4.4	Glitch statistics	100
4.4.5	New and updated pulsar parameters	101
4.5	Braking indices and long-term evolution	104
4.5.1	Observed versus predicted evolution in $P-\dot{P}$ space	104
4.5.2	Connecting large braking indices to glitches in Vela-like pulsars	111
4.5.3	Pulsars with seemingly small braking indices	116
4.6	Conclusions	118
5	Swift J1818.0–1607: A 1.4 s radio magnetar	121
5.1	Introduction	121
5.2	Observation and analysis	123
5.2.1	Profile phenomenology and flux density	125
5.2.2	Polarimetry	126
5.2.3	Single pulses	127
5.3	Discussion	131
6	The dynamic magnetosphere of Swift J1818.0–1607	135
6.1	Introduction	136
6.2	Observations	138
6.3	Profile and spectral evolution	139
6.4	Emission mode switching	143
6.5	Polarization properties and geometry	147
6.5.1	Viewing geometry	148
6.5.2	Position angle alignment and emission heights	152
6.5.3	Polarization modes	154
6.6	Discussion	156
6.7	Summary and conclusion	162
7	Shining a light through a pulsar magnetosphere	165
7.1	Introduction	165
7.2	MeerKAT observations	167
7.3	Eclipse light-curve modelling	167
7.3.1	The Lyutikov and Thompson light-curve model	168
7.3.2	Joint-fitting of eclipse pairs	172
7.3.3	Spin-orbit precession of pulsar B	174
7.4	Spectropolarimetry of the eclipses	179

7.4.1	Polarised light curves	179
7.4.2	Position angle variations and Faraday rotation	181
7.4.3	Birefringence induced circular polarisation	183
7.5	Conclusions	188
8	Conclusions and future prospects	191
8.1	Glitches, timing noise and long-term rotational evolution	191
8.2	Swift J1818.0–1607 and other radio-loud magnetars	195
8.3	The Double Pulsar	198
8.4	Final remarks	201
	Bibliography	225
A	UTMOST timing results	227
B	Hidden-Markov Model glitch detector	241
B.1	HMM recipe and parameters	241
B.2	Glitch upper-limits	243
C	Parkes timing results	245

List of Figures

1.1	Lighthouse model of pulsars.	2
1.2	Polarisation vector on the Poincaré sphere	6
1.3	Diagram of the rotating vector model.	7
1.4	Effects of birefringence on the Poincaré sphere.	12
1.5	$P-\dot{P}$ diagram of catalogued pulsars.	14
1.6	Pulsar formation pathways.	16
2.1	Basic signal path of a radio receiver	28
2.2	The CSIRO Parkes 64-m radio telescope.	31
2.3	East-arm of the Molonglo Observatory Synthesis Telescope.	34
2.4	Antennas in the MeerKAT core.	35
2.5	Diagram showing the pulsar timing process.	37
2.6	Schematic view of a Keplerian orbit.	40
3.1	$P-\dot{P}$ diagram of pulsars monitored by UTMOST	49
3.2	Timing residuals after fitting for ν and $\dot{\nu}$	56
3.3	Timing residuals of four glitched pulsars.	62
3.4	Post-fit timing residuals of PSR J1825–0935.	66
3.5	Posterior distributions for the recovered scaling relation.	74
3.6	Comparison of scaling relation coefficients with values from the literature.	77
3.7	Average pulse profile of PSR J0737–3039A at 835 MHz.	78
3.8	Sub-integration stack for PSR J1402–5124 at 835 MHz	80
4.1	Corner plot for glitch in PSR J1524–5625.	91
4.2	Histograms of glitch-induced changes in ν and $\dot{\nu}$	96
4.3	Histogram of fractional glitch recovery parameters.	100
4.4	Posterior distributions for PSR J0908–4913.	103
4.5	Braking indices measured in Parkes young pulsars.	106
4.6	Observed long-term rotational evolution in $P-\dot{P}$ space.	107
4.7	Posterior distributions for PSR J1734–3333.	111
4.8	Inter-glitch $\ddot{\nu}$ versus $\Delta\dot{\nu}_g/T_g$	115
4.9	$P-\dot{P}$ diagram of young pulsars monitored by Parkes.	117
5.1	Swift J1818.0–1607 polarisation spectra and integrated profiles.	123
5.2	Sub-banded averaged flux densities and scattering timescales.	126

5.3	Sub-banded polarisation profiles of Swift J1818.0–1607	128
5.4	Matched-filter single-pulse S/N distribution of Swift J1818.0–1607	129
5.5	A single pulse from Swift J1818.0–1607.	130
6.1	Parkes UWL observations of Swift J1818.0–1607.	140
6.2	Mode-switching detected on MJD 59009.	144
6.3	Mode-switching detected on MJD 59047.	145
6.4	Continuum flux densities for the P- and M-modes.	146
6.5	Position angle swings of Swift J1818.0–1607 with RVM fits overlaid.	149
6.6	Posterior distributions of α and ζ	150
6.7	Posterior distributions of α and β	151
6.8	PA swings and total intensity profiles aligned to MJD 59087 RVM fit.	152
6.9	Pulse and phase resolved polarization position angle.	154
6.10	Poincaré sphere projection of Ψ and χ measurements from MJD 59109.	156
6.11	Variations in the spin-frequency of Swift J1818.0–1607 over time.	159
7.1	Diagram of the eclipse light-curve model.	169
7.2	Eclipse light curves detected using the L-band and UHF receivers.	174
7.3	Evolution of pulsar B’s geometry over time.	175
7.4	Posterior distributions for the geodetic precession rate.	177
7.5	Mass-mas diagram for pulsars A and B.	178
7.6	Polarisation profile of pulsar A at L-band and UHF.	180
7.7	Orbital-phase averaged polarisation light curve across the eclipse.	182
7.8	Stokes spectra at the ingress Stokes V peak.	185
7.9	Stokes spectra at the superior conjunction Stokes V peak.	186
7.10	Stokes spectra at the egress Stokes V peak.	187
7.11	Ascribing detected polarisation variations to pulsar B’s geometry.	188

List of Tables

3.1	Priors on pulsar and timing noise parameters.	51
3.2	Maximum likelihood values for $\dot{\nu}$ and n	60
3.3	Maximum likelihood glitch parameters.	64
3.4	Recovered glitch parameters for PSR J1825–0935	67
3.5	Correlations between various pulsar properties and red noise strength.	71
3.6	Comparison between recovered scaling relation parameters.	75
4.1	Priors for intrinsic, stochastic and glitch parameters.	89
4.2	Recovered glitch properties.	93
4.3	Pulsars that display significant proper-motions.	102
4.4	Braking index measurements.	105
4.5	Inter-glitch braking index measurements.	113
5.1	Measurements of pulse scattering, flux and polarisation fraction.	124
6.1	List of Parkes UWL observations of Swift J1818.0–1607.	139
6.2	Results from spectral fits and associated log Bayes factors.	142
6.3	Results from rotating vector model fits.	148
6.4	Longitude offsets and relative emission heights.	153
7.1	Priors for the light-curve model parameters.	173
A.1	Astrometric and rotational properties inferred from UTMOST timing.	227
A.2	Preferred timing model and stochastic red noise properties.	233
B.1	Pulsar averaged 90% upper-limits on $\Delta\nu_g/\nu$	244
C.1	Preferred timing models and associated Bayes factors.	245
C.2	Astrometric and rotational parameters for 51 glitched pulsars.	246
C.3	Astrometric and rotational parameters for 23 non-glitched pulsars.	248

1

Introduction

1.1 Neutron stars and pulsars

Neutron stars are dense, compact objects with masses ranging between 1.1 - $2.1 M_{\odot}$ and radii of only ~ 10 km, making them some of the densest objects in the Universe. They are one of the end products of massive stellar evolution, and thought to be born from the core-collapse supernovae of stars initially 10 - $25 M_{\odot}$ (Kaspi & Helfand, 2002). Most known neutron stars are observed as pulsars, a type of highly magnetised, rotating neutron star that are detected via (apparent) pulsations that are locked to the rotation of the star.

While the precise details of the pulsar emission process are still up for debate (Melrose & Rafat, 2017), it is generally thought the rapid rotation of the neutron star, along with its co-rotating magnetic field, creates an electric field that draws charged particles from the neutron star surface into a plasma contained within the magnetosphere. Particles streaming along closed magnetic field lines will co-rotate with the neutron star, whereas those on open magnetic field lines near the poles will be accelerated away from the neutron star and produce highly beamed, coherent radiation (Pacini, 1968; Gold, 1969; Goldreich & Julian, 1969). If these beams are misaligned to the pulsar rotation axis, then the beams will sweep out across our Galaxy, where those that cross our line of sight on Earth may be detectable as sources of pulsed electromagnetic radiation, typically at radio wavelengths between ~ 100 MHz to 10 GHz. A cartoon of the canonical ‘lighthouse’ model of radio pulsars, alongside a simplified diagram of their internal structure, is shown in Figure 1.1. High-energy emission (~ 0.1 keV to 100 MeV) is also detected from a significant fraction of pulsars. Unlike the radio emission, the gamma-ray and X-ray pulses are believed to originate from the acceleration of charged particles near the pulsar light-cylinder (Harding et al., 1978; Romani, 1996; Abdo et al., 2013) and thermal hotspots on the neutron star surface where charged particles flow down open magnetic field-lines and collide with the

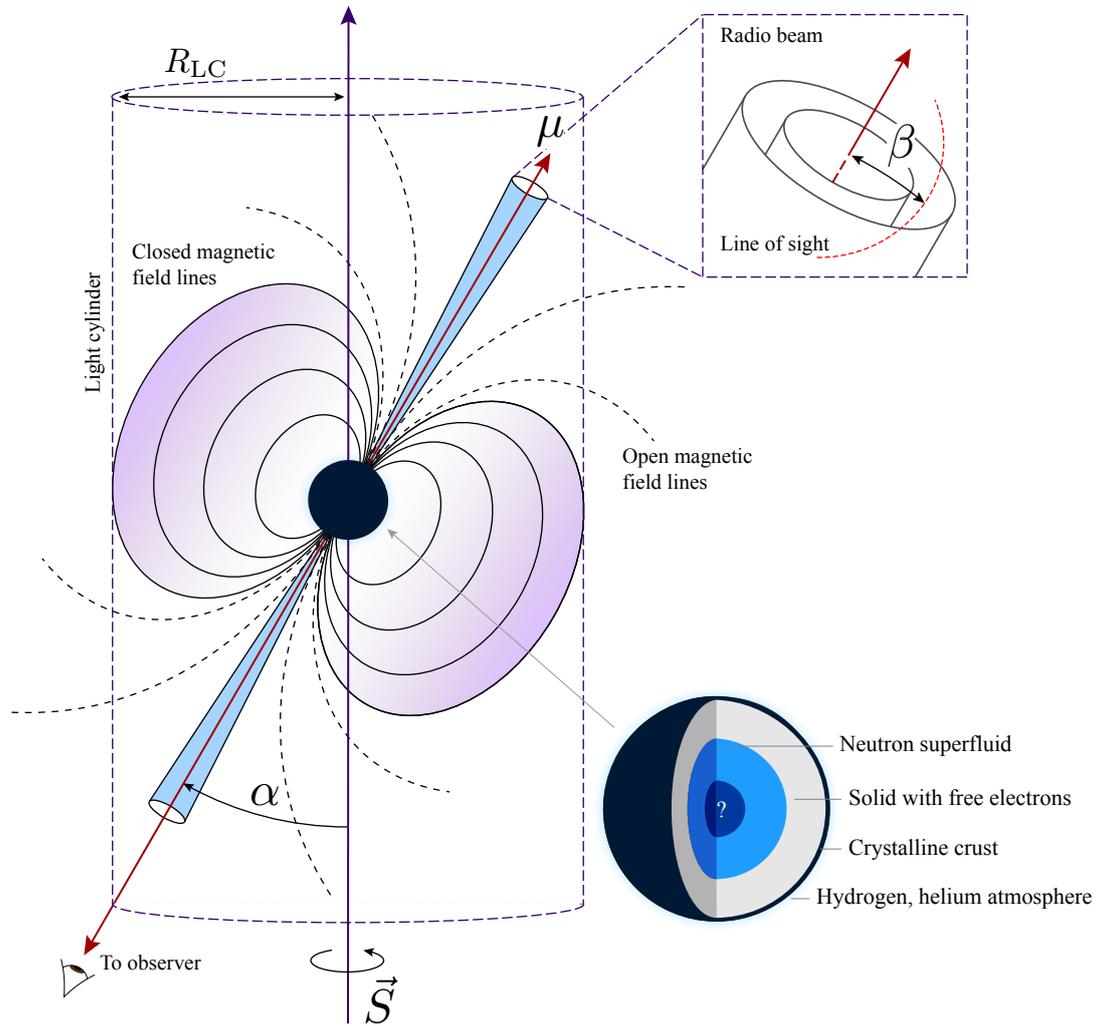


Figure 1.1 Diagram of the lighthouse model of pulsars. The neutron star (dark blue) rotates about a spin-vector (\vec{S}) with a magnetic axis (μ) that has an angular offset (α) from the spin axis. The closed, dipole magnetosphere of the pulsar extends out to the light-cylinder radius (R_{LC}), where the co-rotating velocity of exceeds the speed of light. The inset image in the upper-right depicts the sweep of the emission cone across our line of sight, offset from the magnetic pole (by an angle β). A cross-section of the potential interior structure of a neutron star is shown in the bottom right.

outer crust (Wang et al., 1998).

Since their discovery ~ 50 years ago by Jocelyn Bell-Burnell (Hewish et al., 1968), the study of pulsars has enabled astronomers to explore a wide variety of fundamental physics that are otherwise inaccessible to experiments located here on Earth or within the Solar System. This includes probing the behaviour of baryonic matter under super-nuclear densities (Özel & Freire, 2016), studying electrodynamics and plasma physics in extreme

magneto-ionic environments (Eatough et al., 2013) and conducting some of the most stringent tests of Einstein’s theory of general relativity in the strong-field regime (Kramer et al., 2006b). They are also useful as tools to probe the gas and magnetic fields within the ionised interstellar medium (ISM) of the Milky-Way (e.g. Cordes, 1986) and as low-frequency gravitational-wave detectors in pulsar timing arrays (Detweiler, 1979; Hellings & Downs, 1983).

1.1.1 Internal structure

As depicted in Figure 1.1, the interiors of neutron stars are believed to be differentiated into a series of layers surrounded by a diffuse hydrogen and helium atmosphere (Lattimer & Prakash, 2004; Watts et al., 2015). The solid outer crust is comprised of atomic nuclei, predominately iron and other heavy elements, arranged into a crystalline lattice. This transitions into a solid layer of neutron-rich atomic nuclei within which electrons are able to freely flow. Deeper still is the outer core, where neutrons have dripped out into a superfluid state (alongside a non-negligible fraction of superconducting electrons and positrons) that is threaded by vortices attached to pinning regions within the inner crust. Interactions within the outer core are dominated by nuclear forces as opposed to electromagnetism in the crust. Finally, at the centre of a neutron star is the inner core. The exact state and behaviour of baryonic matter at super-nuclear densities within the inner core is unclear. Various neutron star ‘equations of state’ (EoS) predict a variety of phase-transitions within the inner core that may allow for the stable existence of exotic states of matter, such as deconfined quark matter, strangelets, and condensates of pions or kaons (Lattimer, 2012). Different equations of state yield distinct neutron star mass-radius relationships, among other potentially observable properties (Lattimer & Prakash, 2004; Özel & Freire, 2016)

Our constraints on the nuclear EoS have been vastly improved over the previous decade thanks to several major breakthroughs. These included the discovery of ultra-massive neutron stars such as PSRs J1614–2230 (Demorest et al., 2010) and J0740+6620 (Cromartie et al., 2020), both of which reside near the maximum possible neutron star mass. Additional constraints came from both the first observation of gravitational waves from a double neutron star merger, GW170817 (Abbott et al., 2017a,b), which enabled us to constrain the tidal deformability of neutron stars (Abbott et al., 2018), and measurements of neutron star radii via relativistic ray-tracing of pulsar X-ray pulsations using the Neutron Star Interior Composition Explorer (NICER; Miller et al. 2019; Riley et al. 2019, 2021). Combinations of these results along with measurements of the neutron skin of ^{208}Pb from the PREX-2 experiment (Adhikari et al., 2021; Reed et al., 2021), show that neutron star

radii do not vary significantly with increasing mass, thereby implying a stiff nuclear EoS is preferred over comparatively soft EoS models (Raaijmakers et al., 2021; Pang et al., 2021).

1.1.2 General properties

Approximating a pulsar as a solid, rigidly rotating sphere, its spin-period, P , increases slowly over time at a rate \dot{P} as it radiates away angular momentum. The energy loss-rate of a pulsar, also referred to as its spin-down luminosity (\dot{E}), can be defined as

$$\dot{E} = -I\Omega\dot{\Omega} = 4\pi^2 I \frac{\dot{P}}{P^3}, \quad (1.1)$$

where I is the neutron star moment of inertia and $\Omega = 2\pi/P$ is the angular rotation frequency of the pulsar. The precise value of I depends on the mass and radius of a neutron star, which are in turn related to the unknown neutron star equation of state. For a canonical neutron-star with a mass of $1.4 M_\odot$ and 10 km radius, $I \approx 10^{45} \text{ g cm}^2$. Under the assumption that a pulsar can be approximated to a rotating dipole in a vacuum, the rotational energy lost due to dipole radiation is simply (Jackson, 1998)

$$\dot{E}_{\text{dipole}} = \frac{2|\mathbf{m}|^2 \sin^2 \alpha}{3c^3} \Omega^4 \quad (1.2)$$

where $|\mathbf{m}|$ is the magnetic dipole moment. Equating Equations 1.1 and 1.2 we obtain

$$\dot{\Omega} = -\left(\frac{2|\mathbf{m}|^2 \sin^2 \alpha}{3Ic^3}\right) \Omega^3, \quad (1.3)$$

which can be further generalised such that (Gunn & Ostriker, 1969)

$$\dot{\Omega} = -\kappa\Omega^n \Leftrightarrow \dot{P} = -\kappa P^{2-n} \quad (1.4)$$

where κ is positive, time-independent constant and n is the braking index of the pulsar.

The age of a pulsar can be inferred by integrating Equation 1.4 and substituting in $\kappa = \dot{P}/P^{2-n}$

$$\tau = t - t_0 = \int_{P_0}^P \frac{1}{\kappa P^{2-n}} dP = \frac{P}{\dot{P}(n-1)} \left[1 - \left(\frac{P_0}{P}\right)^{n-1}\right], \quad (1.5)$$

where t_0 is the time at which the pulsar was born and P_0 its initial spin-period. Assuming the pulsar evolves purely due to dipole radiation ($n = 3$) and the currently observed spin-

period is much larger than it was at birth ($P \gg P_0$), then Equation 1.5 simplifies to the ‘characteristic age’

$$\tau_c = \frac{P}{2\dot{P}}. \quad (1.6)$$

Characteristic ages have occasionally been shown to not reflect the true age of a pulsar but nevertheless provide a reasonable estimate of how old they may be.

In Equation 1.3, the dipole magnetic moment can be written in terms of the neutron-star radius (R_{NS}) and magnetic field strength as $|\mathbf{m}| = BR_{\text{NS}}^3$ (Jackson, 1998). Assuming the neutron star is an orthogonal rotator ($\alpha = 90^\circ$) with a radius of 10 km and $I \approx 10^{45} \text{ g cm}^2$, then the magnetic field strength at its surface is simply

$$B_{\text{surf}} = \sqrt{\frac{3Ic^3}{8\pi^2}} \frac{P\dot{P}}{R^6 \sin^2 \alpha} \simeq 3.2 \times 10^{19} \text{ G} \sqrt{P\dot{P}}. \quad (1.7)$$

Like characteristic ages, this relation returns a simplified approximation of true magnetic field strength. Near their surfaces, many pulsars are believed to have significant multipolar components to their magnetic fields (Arons, 1993; Asseo & Khechinashvili, 2002), hence the values returned from Equation 1.7 can in some sense be considered lower-limits on B_{surf} .

1.1.3 Polarimetry

The detection of strongly polarised emission from the Vela pulsar (PSR J0835–4510) by Radhakrishnan & Cooke (1969) demonstrated that pulsars are strong, persistent sources of polarised radiation. Polarised radiation can be described in terms of the Stokes parameters $\mathbf{S} = \{I, Q, U, V\}$, where I is the total intensity, Q and U are the two orthogonal linear polarisation components, the intensity of which is given by $L = \sqrt{Q^2 + U^2}$, and V is the circular polarisation. Under the PSR/IEEE convention, circular polarisation with a positive sign is right-handed and a negative sign is left-handed (van Straten et al., 2010).

The Stokes parameters can be related to the observed electromagnetic field detected by two orthogonal linear receivers as the sum of its two orthogonal components E_x and E_y as

$$\begin{aligned} I &= \langle E_x^2 \rangle + \langle E_y^2 \rangle, \\ Q &= \langle E_x^2 \rangle - \langle E_y^2 \rangle, \\ U &= 2\langle E_x \rangle \langle E_y \rangle \cos(\theta), \\ V &= 2\langle E_x \rangle \langle E_y \rangle \sin(\theta), \end{aligned} \quad (1.8)$$

where ‘ $\langle \dots \rangle$ ’ indicates a time-averaging and θ is the phase difference between the two

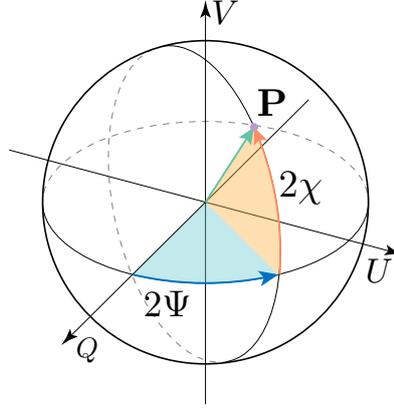


Figure 1.2 Polarisation vector (\mathbf{P}) projected onto the Poincaré sphere with a co-latitude and co-longitude of 2χ and 2Ψ respectively.

orthogonal components. For an electromagnetic wave that is 100% polarised, the total intensity is simply the quadrature sum of the the polarisation parameters, $I^2 = Q^2 + U^2 + V^2$. However, not all radiation is fully polarised, and may instead have a finite polarisation fraction, P/I , where $P = \sqrt{Q^2 + U^2 + V^2}$ is the total polarisation. In spherical polar coordinates, the Stokes Q , U and V components can be rewritten as

$$\begin{aligned} Q &= P \cos(2\Psi) \cos(2\chi), \\ U &= P \sin(2\Psi) \cos(2\chi), \\ V &= P \sin(2\chi), \end{aligned} \tag{1.9}$$

where $\Psi = 0.5 \tan^{-1}(U/Q)$ is the linear polarisation position angle (PA) and $\chi = 0.5 \tan^{-1}(V/L)$ is the ellipticity angle. It is often useful to project the Stokes Q , U and V components, written in terms of the normalised polarisation vector as

$$\hat{\mathbf{P}} = \frac{1}{P} \begin{bmatrix} Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} \cos(2\Psi) \cos(2\chi) \\ \sin(2\Psi) \cos(2\chi) \\ \sin(2\chi) \end{bmatrix}, \tag{1.10}$$

onto the Poincaré sphere, as depicted in Figure 1.2.

[Radhakrishnan et al. \(1969\)](#) suggested the sweep of the PA across the pulse profile of the Vela pulsar is a purely geometric effect induced by the changing angle between the sky-projected dipole magnetic field of the pulsar and our line of sight. This ‘rotating vector model’ (RVM) assumes the radio emission is tied to the field lines near the magnetic pole(s)

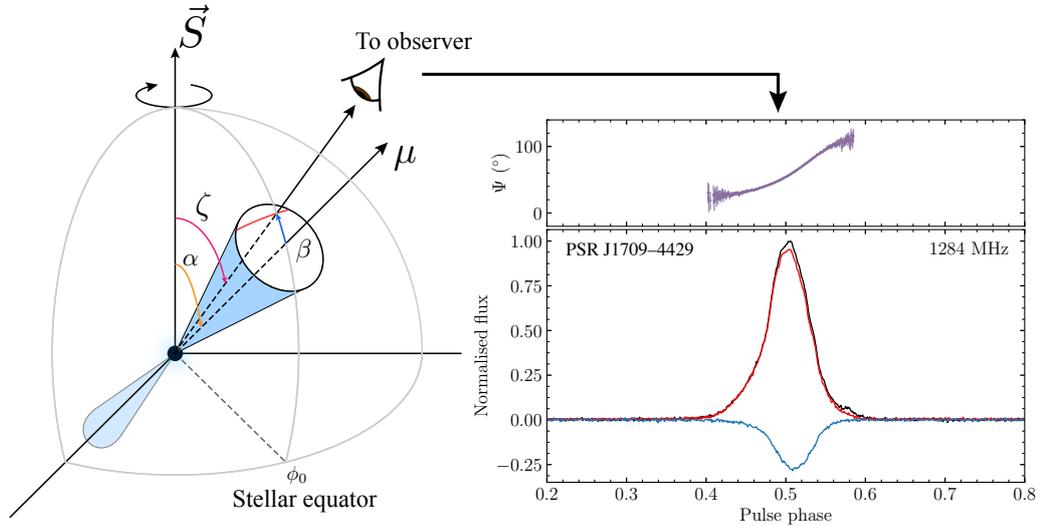


Figure 1.3 Example of the position angle (Ψ) sweep seen in PSR J1709–4429 by MeerKAT. The left-hand side depicts a cartoon of the rotating vector model of the pulsar, where our line of sight is shown by the red trace across the emission cone. The right-hand side shows the resulting polarisation profile and position angle swing, where the black curve indicates total intensity, red corresponds to linear polarisation and blue is circular polarisation.

of the neutron star and is emitted perpendicular to the local magnetic-field direction. The RVM predicts the PA swing should follow S-shaped curve as a function of pulse phase or longitude (ϕ) as

$$\tan(\Psi - \Psi_0) = \frac{\sin \alpha \sin(\phi - \phi_0)}{\sin \zeta \cos \alpha - \cos \zeta \sin \alpha \cos(\phi - \phi_0)}, \quad (1.11)$$

where α is the magnetic inclination angle, ζ is the angle between the pulsar spin-axis and our line of sight ($\beta = \zeta - \alpha$ is the impact angle between the magnetic axis and our line of sight), and Ψ_0 is the central PA at the fiducial longitude ϕ_0 . An example of this is shown for the young pulsar PSR J1709–4429 in Figure 1.3. RVM fits to pulsar PA swings enables us to constrain both the magnetic geometry of pulsars, as well as their viewing angles (e.g. Johnston & Kramer, 2019). It can also be used to measure the lag between the total intensity profile and the PA swing that arises from relativistic aberration and retardation of the pulsar emission which can be used to infer the radio emission height above the neutron star surface (Blaskiewicz et al., 1991). Studies that compare offsets between the inflection point of the PA and the total intensity profile across a wide range of pulsars indicate the radio emission originates from between ~ 10 -1000 km above the

neutron star surface (Blaskiewicz et al., 1991; Mitra & Li, 2004; Johnston & Karastergiou, 2019; Johnston & Kramer, 2019; Desvignes et al., 2019).

While the RVM is often considered the standard model for describing pulsar polarimetry, there are many pulsars that deviate significantly away from the predicted S-shaped PA swing (e.g. Johnston & Weisberg, 2006; Johnston et al., 2008; Johnston & Kerr, 2018). Such phenomena can arise from the propagation effects within the pulsar magnetosphere (e.g. Ilie et al., 2019) or the presence of multiple, orthogonally polarised emission modes (Manchester, 1975; Karastergiou et al., 2002).

In general, the structure of the magnetospheres of pulsars are thought to be largely static, with any secular evolution taking place over timescales much greater than a human lifespan (e.g magnetic-field decay; Ho et al., 2015). However, there is a growing population of so-called ‘mode-switching’ pulsars that are observed to stochastically switch between two or more quasi-stable emission states (Backer, 1970b,a; Lyne et al., 2010) over timescales as short as a few rotations. Correlations between the switching and spin-down rate of these pulsars suggests it results from variations in the plasma outflows associated with the radio emission mechanism (Kramer et al., 2006a; Lyne et al., 2010). There are also a handful of extremely magnetised pulsars that possess dynamic magnetospheres, where the profile shape and polarisation undergo dramatic variations over timescales as short as a few minutes to months. One of these objects is the focus of Chapters 5 and 6.

1.2 Propagation effects

As the radio emission from a pulsar travels toward the Earth, it passes through the plasma that makes up the ISM along the line of sight. This introduces a variety of propagation effects that can be used to infer the physical properties of the ISM itself and structures within the Milky-Way.

1.2.1 Dispersion

Electrostatic interactions between the radio waves emitted by a pulsar and free-electrons within the ISM introduces a frequency dependent delay (dispersion) in their arrival time at the Earth. Radio waves with a frequency (f) that is greater than the plasma frequency of the intervening electrons, i.e $f > f_p \simeq 8.5 \text{ kHz} \sqrt{n_e/\text{cm}^{-3}}$, will undergo a reduction in their group velocity $v_g(f) = c\mu(f)$, where c the vacuum speed of light and $\mu(f) = \sqrt{1 - (f_p/f)^2}$ is the refractive index. This effect was noticed in the first pulsar discovery by Hewish et al. (1968), where the higher-frequency radio waves were seen to arrive earlier than

those at lower frequencies. The resulting time delay experienced by a radio pulse at a given observing frequency is given by

$$\Delta t_{\text{DM}} = \left(\int \frac{d\ell}{v_g(f)} \right) - \frac{L}{c} = \mathcal{D} \int \frac{n_e}{f^2} d\ell, \quad (1.12)$$

where L is the distance to the pulsar, $\mathcal{D} = 4.15 \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s}$ is the dispersion constant, ℓ is the path length between the Earth and the pulsar and n_e is the total electron column density. An expression for the time delay between two discrete observing frequencies, f_1 and f_2 (in units of MHz), can be derived from Equation 1.12 as

$$\Delta t_{\text{DM}} = \mathcal{D} (f_1^{-2} - f_2^{-2}) \text{DM}. \quad (1.13)$$

Here, DM is the ‘dispersion measure’ of the pulse, which encodes the line-integral of the total electron column density along the line of sight in units of pc cm^{-3} (Lorimer & Kramer, 2012)

$$\text{DM} = \mathcal{D} \int n_e d\ell. \quad (1.14)$$

Strictly speaking the DM constant is only roughly proportional to the number of free electrons, as heavier ions also contribute to the dispersion.

DM measurements can be used in conjunction with models of the Galactic free-electron density to infer the distances to pulsars, albeit with uncertainties that can be as high as 50-90%. The first such electron density model was published by Manchester & Taylor (1981), followed by the LMT85 model (Lyne et al., 1985). The more complicated TC93 model (Taylor & Cordes, 1993) attempted to account for density variations due to the spiral-arms of the Milky-Way, as well as pulse scatter broadening. The TC93 model has since been superseded by the NE2001 model (Cordes & Lazio, 2002), which has been the standard electron density model in pulsar astronomy for nearly two decades. More recently, the YMW16 model (Yao et al., 2017) was developed in an attempt to account for a number of nearby structures in the ISM.

1.2.2 Scattering and scintillation

Turbulent regions within the ISM can act on radio waves passing through it in lens-like fashion, inducing time-dependent focussing, refocusing and steering of the transmitted radiation. This results in observable effects such as refractive and diffractive scintillation, and frequency dependent scattering of the detected pulse profile (Cordes et al., 1986). If we approximate the ISM as a thin screen located along the line-of-sight between us and

the pulsar (Williamson, 1972), then the timescales over which refractive and diffractive scintillation vary can be used to infer the relative transverse velocity of the pulsar, ISM and the Earth (Cordes, 1986; Cordes & Rickett, 1998), this can in turn enable accurate distance measurements to the pulsar. In some cases it is possible to use pulsar scintillation to probe turbulent and discrete structures in the ISM plasma (Cordes & Downs, 1985; Stinebring et al., 2000; Brisken et al., 2010). It can also be used to measure the geometry of pulsar binaries independent of pulsar timing, thereby enabling improved tests of GR. An example of this analysis technique applied to the relativistic pulsar-white dwarf binary PSR J1141–6545 are presented in works by Ord et al. (2002) and Reardon et al. (2019).

Multi-path propagation due to interstellar scattering also causes the observed pulse profile to become smeared when observed at low radio frequencies. This effect, known as pulse scatter broadening, is often modelled by convolving the pulse profile with a one-sided exponential pulse broadening function as

$$O(t) = S(t) \otimes e^{-t/\tau_{sc}}, \quad (1.15)$$

where $S(t)$ is the un-scattered pulse profile, \otimes indicates a convolution, and τ_{sc} is the scattering timescale. Studies of how pulsar scattering varies as a function of frequency in large samples of pulsars have shown it is often well approximated by a power-law of the form $\tau_{sc}(f) \approx f^{\alpha_{sc}}$ with measured scattering indices (α_{sc}) often found to be between -4 to -4.4 (e.g. Geyer et al., 2017; Oswald et al., 2021). These measurements indicate the intervening scattering screen is usually dominated by Kolomogorov turbulence (Geyer & Karastergiou, 2016).

1.2.3 Faraday rotation

If the ISM is magnetised and has a magnetic field component that runs parallel to the direction in which the radio waves propagate, then the linearly polarised component will undergo a rotation about the polarisation plane (Smith, 1968b). This effect, known as Faraday rotation, is due to the magnetised ISM having a polarisation-dependent refractive index that causes right- and left-handed circularly polarised waves to propagate at different speeds. Faraday rotation introduces a phase shift in the observed polarisation position angle that depends on the integrated magnetic field strength and electron column density along the line of sight. This results in a wavelength dependent rotation of the polarisation vector about the V-axis at a fixed latitude on the Poincaré sphere, as depicted in the left-hand side of Figure 1.4. An expression describing the effect can be drawn from Equation

1.10 as

$$\hat{\mathbf{P}}(\lambda) = \begin{bmatrix} \cos[2\Psi(\lambda)] \cos(2\chi) \\ \sin[2\Psi(\lambda)] \cos(2\chi) \\ \sin(2\chi) \end{bmatrix}, \quad (1.16)$$

where the wavelength dependence of the PA is given by

$$\Psi(\lambda) = \Psi_0 + \text{RM}(\lambda^2 - \lambda_c^2). \quad (1.17)$$

Here, Ψ_0 is the PA at some reference wavelength λ_c (often chosen to be the centre frequency of the input spectrum) and RM is the rotation measure. The RM is dependent on the line-of-sight electron column density and parallel magnetic field strength (B_{\parallel}) as

$$\text{RM} = \frac{e^2}{8\pi^2 \varepsilon_0 m_e^2 c^3} \int dl n_e B_{\parallel}, \quad (1.18)$$

where e is the electron charge, ε is the vacuum permittivity, and m_e the electron rest mass. Whenever the DM and RM of a pulsar are both known, the average parallel magnetic-field strength (in micro-Gauss; μG) along the line of sight to the pulsar can be inferred via the relation (Smith, 1968b,a; Ekers et al., 1969)

$$\langle B_{\parallel} \rangle = \frac{\int n_e B_{\parallel} dl}{\int n_e dl} = 1.23 \frac{\text{RM}}{\text{DM}} \mu\text{G}, \quad (1.19)$$

where RM and DM are in their natural units of rad m^{-2} and pc cm^{-3} respectively. Polarisation studies of large samples of pulsars and quasars have enabled mapping of the strength and direction of magnetic fields within the Milky-Way spiral arms and bulge (see recent studies by Han et al., 2018; Sobey et al., 2019), in addition to magnetic fields of the Magellanic Clouds (Gaensler et al., 2005; Mao et al., 2008) and some globular clusters (e.g. Abbate et al., 2020).

A generalised form of Faraday rotation (sometimes referred to as Faraday conversion) occurs when the natural wave modes of the propagating medium are either linearly or elliptically polarised, resulting in a conversion between the linear and circular polarisation components of the incident radiation (Sazonov, 1969; Pacholczyk & Swihart, 1970; Kennett & Melrose, 1998). This phenomenon can be visualised as a rotation of the polarisation vector on the Poincaré sphere about an axis tilted away from the Stokes V pole, as depicted by the middle and right-hand side of Figure 1.4. Note, the angular offset between the axis of rotation and the Stokes V pole will reflect the natural wave-modes of the intervening medium the radiation has propagated through and can hence be used to assess

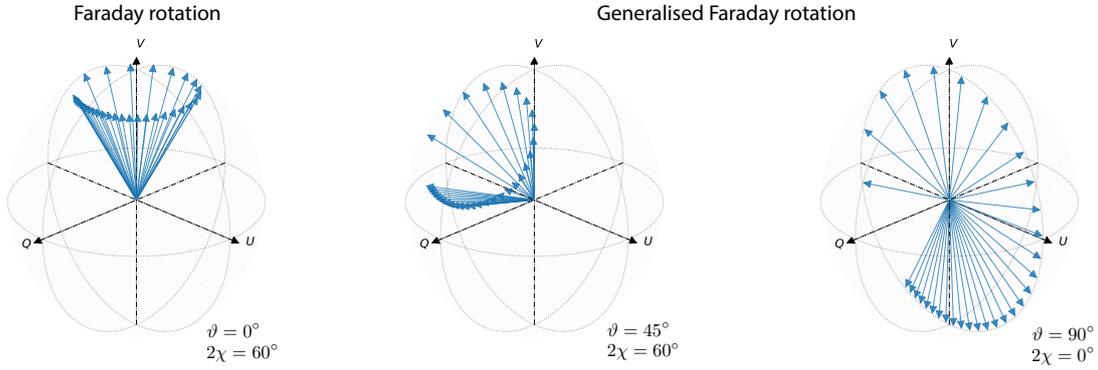


Figure 1.4 Polarization vector in terms of Stokes Q , U and V projected onto the Poincaré sphere, demonstrating the effects of standard Faraday rotation (left) and generalised Faraday rotation (middle and right).

the properties of the medium. For instance, a relativistic plasma comprised of electron-positron pairs will have linearly polarised wave-modes, resulting in an axis of rotation that would point into the Stokes Q - U plane (e.g. [Sazonov 1969](#)). The tilting of the polarisation plane can be emulated through the addition of two rotation matrices to Equation 1.16,

$$\mathbf{R}_\vartheta = \begin{bmatrix} \cos(\vartheta) & 0 & \sin(\vartheta) \\ 0 & 1 & 0 \\ -\sin(\vartheta) & 0 & \cos(\vartheta) \end{bmatrix}, \quad (1.20)$$

and

$$\mathbf{R}_\varphi = \begin{bmatrix} \cos(\varphi) & -\sin(\varphi) & 0 \\ \sin(\varphi) & \cos(\varphi) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (1.21)$$

where the angles ϑ and φ represent respective rotations about the Stokes U and V axes. Hence, a full phenomenological generalised Faraday rotation model can be written as

$$\hat{\mathbf{P}}_m(\lambda) = \mathbf{R}_\vartheta \cdot \mathbf{R}_\varphi \cdot \mathbf{P}(\lambda). \quad (1.22)$$

The wavelength dependence of Ψ in generalised Faraday rotation can take on different values depending on the underlying physics that governs the propagating medium. Hence, Equation 1.17 is rewritten with an arbitrary wavelength exponent, γ , such that

$$\Psi(\lambda) = \Psi_0 + \text{GRM}(\lambda^\gamma - \lambda_c^\gamma), \quad (1.23)$$

where GRM is the generalised rotation measure with units of $\text{rad m}^{-\gamma}$. Unlike the ‘standard’ RM in Equation 1.17, the functional form of the GRM is not fixed. Instead, it depends on the nature of the intervening medium the polarised radiation has propagated through (e.g. Kennett & Melrose, 1998).

1.3 The neutron star zoo

At the time of writing, there are 3176 known pulsars listed in the Australia Telescope National Facility (ATNF) pulsar catalogue¹ (Manchester et al., 2005). Pulsars are loosely grouped into different categories based on their measured spin period and spin-period derivative. This can be visualised on a $P-\dot{P}$ diagram, such as the one shown in Figure 1.5, where lines of constant characteristic age and surface magnetic field strength (derived via the relations in Section 1.1.2) are overlaid. Four clear groupings of pulsars are visible in the diagram: young pulsars, many of which have associated supernova remnants; the large island of ‘canonical’ pulsars between $P = 10$ ms to 24 s; the smaller millisecond pulsar island with $P < 10$ ms; and the enigmatic magnetars with $B_{\text{surf}} \gtrsim 10^{14}$ G. In this section, we briefly outline how pulsars are expected to evolve through the $P-\dot{P}$ diagram over time, along with the basic properties of the pulsars that fall into each of the aforementioned categories.

1.3.1 Pulsar birth, life and death

Most pulsars are thought to be born towards the upper-left of the $P-\dot{P}$ diagram, with small spin periods ($P \sim 10$ ms), high spin-down rates ($\dot{P} \gtrsim 10^{-15}$), and relatively strong magnetic fields ($B_{\text{surf}} \sim 10^{12}$ G). Many pulsars located within this region of $P-\dot{P}$ space have been kinematically associated with supernova remnants, and have therefore had their young ages confirmed independent of their characteristic spin-down age (see Section 1.1.2). There are also a small number of seemingly young X-ray pulsars that possess spin-down rates that are many orders of magnitude smaller than those of typical young pulsars, yet have unusually high thermal X-ray luminosities. These objects are found to reside in the centres of associated supernova remnants, earning them the moniker ‘central-compact objects’ (CCOs; De Luca, 2017). It is speculated that CCOs are genuinely young pulsars which possess strong magnetic fields that were buried by supernova material falling back onto their surfaces. Their currently buried magnetic fields may re-emerge back to their surfaces via Ohmic diffusion on timescales of 1-100 kyr, transforming them into more

¹<https://www.atnf.csiro.au/research/pulsar/psrcat/>

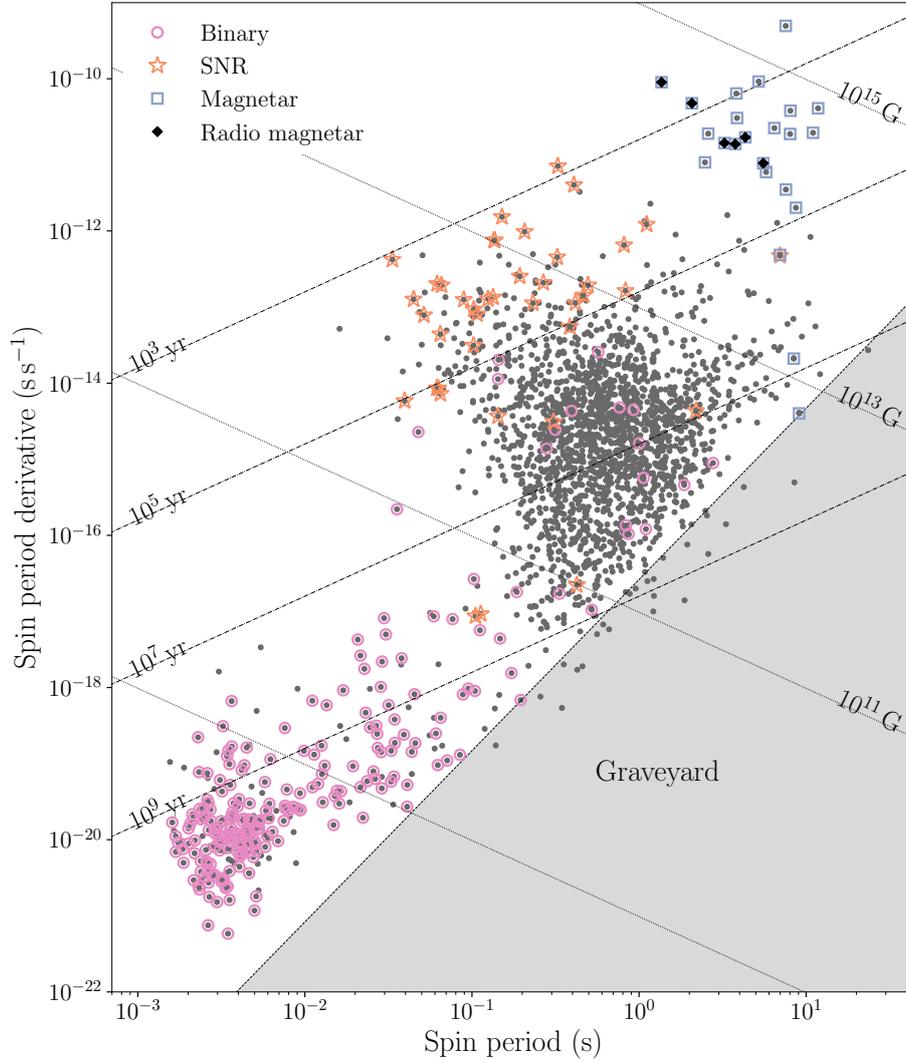


Figure 1.5 Period-period-derivative diagram of the known pulsar population. Pulsars in binary systems are highlighted by magenta circles, those with confirmed supernova (SNR) associations by orange stars, and magnetars by blue squares. Radio-loud magnetars are indicated by black diamonds. Dash-dotted lines indicate lines of constant characteristic age, while dotted lines are for constant magnetic field strength. The radio death line and shaded graveyard assumes the ‘I-prime’ model from Zhang et al. (2000). These data were taken from v1.63 of the ATNF pulsar catalogue (Manchester et al., 2005) using PSRQPY (Pitkin, 2018) and the McGill magnetar catalogue (Olausen & Kaspi, 2014).

‘normal’ young pulsars (Ho, 2011; Viganò & Pons, 2012; Ho et al., 2015). However, the exact long-term evolution of CCOs and how they tie-in to the overall pulsar population remains a mystery.

Studies of young pulsars and their supernova associations assist in solving a range of unsolved issues, including the birth statistics of pulsars (Faucher-Giguère & Kaspi, 2006; Keane & Kramer, 2008; Noutsos et al., 2013), the pulsar emission mechanism, and magnetospheric dynamics through joint radio/high-energy observations (Abdo et al., 2013). Their high spin-down rates mean they are also the only neutron stars for which we can perform studies of long-term pulsar rotational evolution on relatively short timescales ($\sim 10\text{-}50$ yr) (e.g. Hobbs et al., 2004; Livingstone et al., 2007; Lyne et al., 2015). Under the standard picture of pulsar evolution, the rotation period of a pulsar evolves according to the spin-down law in Equation 1.4. If κ and n remain constant over time, they will evolve along straight lines in $P\text{-}\dot{P}$ space, the gradient of which is given by $2 - n$. Pulsars that spin-down purely through dipole radiation (see Section 1.1.2) have $n = 3$, and will therefore follow lines with a slope of -1 (lines of constant magnetic field strength). Note this relatively simplistic picture is only true for pulsars that exist in a vacuum. Particle outflows associated with the pulsar radiation mechanism can aid in carrying angular momentum away from the neutron star into a surrounding pulsar-wind nebula, which can drive pulsar braking indices down to values as low as $n = 1$ (Michel & Tucker, 1969; Melatos, 1997; Contopoulos & Spitkovsky, 2006). Other phenomena that are speculated to occur as neutron stars age over tens of kyr to Myr can induce a time-varying κ term in Equation 1.4, which can result in large observed pulsar braking indices (see recent works by Parthasarathy et al. 2019, 2020). Such processes include the alignment of the magnetic and rotational axes (i.e a decaying α in Figure 1.1; Tauris & Manchester 1998; Johnston & Karastergiou 2017) and magnetic field decay via Hall-drift (Ho et al., 2015). Stochastic irregularities in pulsar rotation rates may also play some role in their long-term evolution, a notion that is explored further in Chapter 4.

As pulsars spin-down, the voltage across their polar caps diminishes. Ruderman & Sutherland (1975) suggested there is a critical pulsar spin period and magnetic field strength above which the electron-position-pair creation driving the pulsar mechanism can continue. This gives rise to the pulsar ‘death line’, below which there is a dearth of pulsars within this region of $P\text{-}\dot{P}$ space, as their spin-down energy is insufficient to support the radio-emission mechanism. Note the positioning of the death line is not only highly model dependent, largely owing to uncertainties surrounding the pulsar emission mechanism (Zhang et al., 2000), but recent discoveries of very-long period pulsars (e.g.

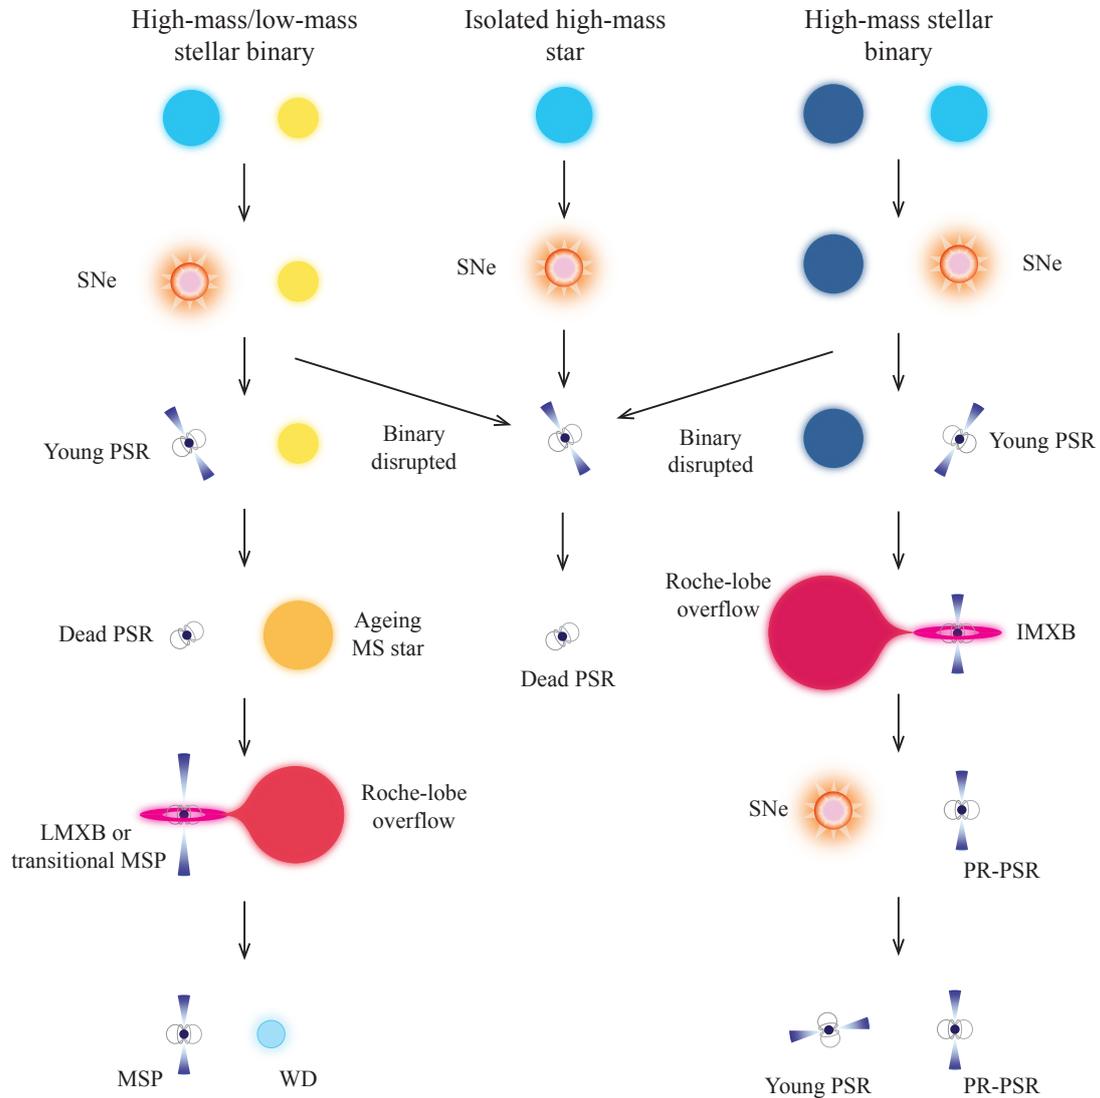


Figure 1.6 Evolutionary paths that result in the formation of pulsar binaries or isolated pulsars.

the 23.5 s PSR J0250+5854; [Tan et al. 2018](#)) suggest reality is much more complex.

1.3.2 Pulsar re-birth via recycling

The second most populous island of pulsars in Figure 1.5 is inhabited by the millisecond pulsars (MSPs). These rapidly rotating neutron stars have spin periods $\lesssim 10$ ms and extremely low spin-down rates ($\dot{P} \lesssim 10^{-18}$). Exactly how they obtain their rapid rotation rates is hinted at by most MSPs being found in binary systems with relatively low-mass, white-dwarf (WD) companions.

Most MSPs are believed to originate from the evolution of stellar binaries that consist of a high-mass O- or B-class star and a low-mass (potentially Sun-like) companion, as summarised in the left-hand side of Figure 1.6. If the supernova of a massive star in a binary system imparts only a small kick on the neutron-star remnant, or if there is insufficient mass-loss for the binary to be disrupted, then the resulting pulsar may be born orbiting a main sequence star. A disrupted binary will result in the pulsar evolving in isolation, as described in the previous section (middle path of Figure 1.6). As the system ages, the neutron star will often follow the standard evolutionary path undertaken by isolated pulsars, passing beyond the death line and ‘dying’ after its radio emission shuts off. Concurrently, the stellar companion will also evolve, eventually exhausting much of its hydrogen fuel, causing it to swell up. If the star fills its Roche lobe, then material from its outer-layers will be pulled into orbit around the neutron star, forming an accretion disc. This transforms the system into a low-mass X-ray binary (LMXB), where friction within the disc causes the material to heat up and radiate at X-ray wavelengths before falling onto the surface of the neutron star, a phenomena that can last up to a few Myr (Alpar et al., 1982; Radhakrishnan & Srinivasan, 1982). The accretion of material onto the neutron star increases its mass, while conservation of angular momentum causes it to spin-up to the point that it rotates with a millisecond spin-period. This pushes it back above the death line in a process termed ‘recycling’ (Bhattacharya & van den Heuvel, 1991). Observational evidence for this process first came from the discovery of millisecond X-ray pulsars by the *Rossi X-ray Timing Explorer* (Wijnands & van der Klis, 1998), as well as ‘transitional millisecond pulsars’ such as PSR J1023+0038 several years later (Bond et al., 2002; Thorstensen & Armstrong, 2005). In the latter systems, an accreting MSP is seen to switch back-and-forth between a radio-silent, accreting LXMB state, and a radio-loud millisecond-pulsar state (Archibald et al., 2009; Stappers et al., 2014). After the companion star completely exhausts its nuclear fuel, it throws off its outer-layers, leaving behind a compact MSP-WD system. The extremely low spin-down rates and incredibly stable spin-periods of MSPs make them some of the most accurate clocks in the Galaxy. Hence, MSPs that reside in binary systems (or even triple systems) with other compact objects make some of the best laboratories for testing our theories of gravity (Stairs, 2003). An ensemble of MSPs distributed across the sky are also actively being monitored by pulsar timing array experiments with the goal of detecting the nanosecond-scale variations induced in their pulse arrival times induced by the stochastic gravitational-wave background (Ferdman et al., 2010; Hobbs et al., 2010; Manchester et al., 2013; McLaughlin, 2013).

Pulsars born with high-mass stellar companions will follow the evolutionary path sum-

marised by the right-hand side of Figure 1.6. These systems often undergo a series of mass-transfer stages (not shown in Figure 1.6 for simplicity), including Roche-lobe overflow from the primary star to the secondary, common-envelope phases that shrink the orbital radius of the system, and intermediate- and high-mass X-ray binary phases (see [Tauris et al. 2017](#) for additional details). As a result of the accelerated evolution of the stellar companion, they do not complete the full recycling process to become MSPs before the companion star also dies in a supernova explosion. If the binary does not become unbound, then the ‘partially recycled’ pulsar (PR-PSR in Figure 1.6) is left in an orbit with another neutron star ([Tauris & van den Heuvel, 2006](#)). These ‘double neutron star’ systems are often much more compact and have higher eccentricities than MSP-WD binaries, making them better laboratories for testing GR in the strong-field regime. Such tests include the first indirect evidence for the existence of gravitational radiation, detected as a shrinking of the orbital period of the Hulse-Taylor binary PSR B1913+16 ([Hulse & Taylor, 1975](#)). Although indirect, this discovery led to the awarding of the 1993 Nobel Prize in physics to Russel Hulse and Joseph Taylor. It also provided the observational basis for the construction of the twin Laser Interferometer Gravitational-wave Observatories, which made the first direct detection of gravitational waves from colliding black holes in September 2015 ([Abbott et al., 2016](#)) and a double-neutron star merger in August 2017 ([Abbott et al., 2017a](#)).

1.3.3 Magnetars

Magnetars are a class of slowly rotating neutron stars characterised by their slow rotation periods and high spin-down rates. As a result, they often have inferred dipole magnetic field strengths in excess of the quantum critical threshold of 4.414×10^{13} G ([Duncan & Thompson, 1992](#); [Paczynski, 1992](#); [Usov, 1992](#)). Unlike ordinary pulsars, their high-energy X-ray luminosity often exceeds their spin-down luminosity, with the excess energy believed to come from the decay of their ultra-strong magnetic fields ([Duncan & Thompson, 1992](#)). To date, 32 magnetars and magnetar candidates have so far been discovered in the Milky-Way and Magellanic Clouds ([Olausen & Kaspi, 2014](#))². A handful of potential magnetar candidates have also been identified in nearby galaxies through the detection of high-energy bursts ([Frederiks et al., 2007](#); [Mazets et al., 2008](#); [Svinkin et al., 2021](#)). Magnetars were first identified as two classes of objects: galactic sources of recurrent gamma-ray and hard X-ray bursts that exhibit structures in their light curves with periodicities ranging between 2-12s (termed soft gamma-ray repeaters; SGRs), and X-ray pulsars that have

²<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>

anomalously high X-ray luminosities (AXPs). Similarities in the rotational properties, quiescent X-ray luminosity, and outburst behaviour of these two types of objects pointed towards a common neutron star source, but with different levels of activity (see [Kaspi & Beloborodov 2017](#) for a general overview). Magnetars spend long periods of time in relative quiescence where very little high-energy emission beyond thermal X-rays are detectable. These relatively inactive phases are interspersed with brief, but intense outbursts often characterised by the emission of bright X-ray and gamma-ray bursts that are believed to be caused by either magnetic re-connection events (e.g. [Lyutikov, 2002](#)) or cracking of the neutron star crust ([Thompson & Duncan, 1995](#)). Timing measurements have shown these outbursts are also associated with large variations in the spin-down rate, likely caused by a combination of particle outflows and variable magnetic torques (e.g. [Scholz et al., 2017](#)).

Radio magnetars

Initial surveys of magnetars with radio telescopes failed to detect pulsed radio emission (e.g. [Lorimer & Xilouris, 2000](#)). [Lyutikov \(2002\)](#) speculated that brief bursts of coherent radio emission may be generated during magnetic re-connection events, with similar phenomenology to radio bursts associated with solar flares. Rapidly fading radio counterparts were detected following high-energy outbursts from SGR 1900+14 ([Frail et al., 1999](#)) and SGR 1806–20 ([Cameron et al., 2005](#)), however they are believed to have resulted from the ejection of ionised clouds as opposed to pulsed radio emission. Other studies predicted pulsar-like coherent emission may be emitted, but the magnetospheric properties of magnetars meant it would be shifted to infrared and optical wavelengths ([Eichler et al., 2002](#)). This radio-silent picture changed with the detection of radio pulses from XTE J1810–197 by [Camilo et al. \(2006\)](#), the first radio-loud magnetar. At the time of writing, only five other magnetars have been found to exhibit pulsed radio emission ([Camilo et al., 2007a](#); [Levin et al., 2010](#); [Shannon & Johnston, 2013](#); [Lower et al., 2020a](#); [Zhang et al., 2020](#)).

The radio emission detected from these objects is transient, only appearing after a high-energy outburst before decaying and eventually disappearing as the magnetar settles into quiescence ([Camilo et al., 2016](#); [Scholz et al., 2017](#)). The emission model proposed by [Beloborodov \(2009\)](#), and later expanded upon by [Wang et al. \(2019\)](#), posits that the radio emission originates from strong electric currents produced by the slow untwisting of bundles of magnetic field lines over time. These field lines are anchored to the rigid neutron star crust, and can become twisted by the sudden crustal motions (i.e a starquake) associated with a magnetar outburst. The size of the emission cone shrinks as the bundle untwists, causing the observed radio profile to narrow over the months to years following

the outburst. Eventually, the magnetar fades from view when the beam edge no longer crosses our line of sight. Such behaviour has been observed in XTE J1810–197 (Camilo et al., 2016) and PSR J1622–4950 (Scholz et al., 2017) following their 2003 and 2011 outbursts.

The post-outburst magnetic fields of magnetars are highly dynamic environments that result in a wide range of polarisation, single-pulse and spectral properties seldom seen among rotation-powered pulsars. Hence, radio-loud magnetars are extremely attractive targets for studying the impact of pulsar magnetospheres on their observed radio phenomenology (Camilo et al., 2007b; Kramer et al., 2007). High time resolution observations of the single pulses from magnetars revealed they are composed of many millisecond-in-width ‘spiky’ sub-pulses with peak flux densities that can exceed ~ 100 Jy (Camilo et al., 2006; Burgay et al., 2018). These short-duration spiky sub-pulses have been found to exhibit striking similarities to the spectral and temporal properties of some fast radio bursts (FRBs), millisecond-duration bursts of radio waves from as-of-yet unknown sources in distant galaxies (see recent reviews by Cordes & Chatterjee 2019 and Petroff et al. 2019, 2021). Recent detections of frequency modulations in radio pulses from SGR 1745–2900 (Pearlman et al., 2018) and XTE J1810–197 (Maan et al., 2019), combined with the numerous FRB progenitor theories that invoke a magnetar central engine (e.g. Wadiasingh & Timokhin, 2019; Margalit et al., 2019; Metzger et al., 2019; Lyutikov & Popov, 2020; Zhang, 2020), suggested a tentative link between magnetars and FRB progenitors. This possible connection was seemingly confirmed by the detection of a bright millisecond-duration radio burst from the galactic magnetar SGR 1935+2154 by the Canadian Hydrogen Intensity Mapping Experiment (CHIME) and Survey for Transient Astronomical Radio Emission 2 (STARE2) projects in North America (CHIME/FRB Collaboration et al., 2020; Bochenek et al., 2020). The 1.5 ± 0.2 MJy-ms fluence of the burst (at 1.4 GHz) means it would have been detectable as an FRB by the 300-m Arecibo telescope were it situated at the distance of the repeating FRB source FRB 180916 (149.0 ± 0.9 Mpc; Marcote et al. 2020). It also suggests that a significant fraction of detected FRBs may originate from similar magnetar-outbursts located at cosmological distances. Hence, studies centred on the Galactic population of radio-loud magnetars may provide an insight into the progenitor objects of extra-galactic FRBs.

1.4 Rotational irregularities in pulsars

Due to their large moments of inertia, pulsars can be used to perform extremely accurate timing experiments by carefully tracking their pulse times of arrival (ToAs). This

technique, referred to as pulsar timing, compares the observed ToAs to a set of predicted arrival times based on a pulsar timing model that encodes all of the known physics of the pulsar (see Chapter 2 for details). However, there are several intrinsic and extrinsic, stochastic processes that can result in the observed ToAs deviating from the predictions of our timing models. Here, we provide an overview of these noise processes and the impact they have on pulsar timing experiments.

1.4.1 White noise and profile jitter

If the observed ToAs from a pulsar are well described by its corresponding timing model, then the resulting timing residuals will exhibit Gaussian or ‘white’ noise that is distributed normally across all fluctuation frequencies. White noise primarily arises from slight variations in the observed profile of a pulsar compared to the template. This can often be attributed to a combination of random fluctuations caused by radiometer noise that arises from the finite sensitivity of the telescope and signal processing backend, as well as intrinsic pulse-to-pulse variations in the profile shape (often termed ‘jitter’; [Groth 1975](#); [Cordes & Downs 1985](#); [Osłowski et al. 2011](#); [Parthasarathy et al. 2021](#)). For pulsars where individual pulses can be detected, jitter can also lead to an underestimation of the formal uncertainty on the ToAs. While radiometer noise can be suppressed by observing pulsars with larger, more sensitive telescopes, eventually the precision to which the pulsar ToAs can be measured will become dominated by jitter. This effectively sets a fundamental limit on how precisely a given pulsar can be timed. Note that more complex timing techniques involving principle-component analysis and direct application of the timing model to pulsar profiles have been able to reduce the impact of jitter in high-precision timing of some millisecond pulsars (see, for example, [Osłowski et al. 2011](#) and [Lentati et al. 2016](#)).

1.4.2 Red timing noise

On timescales of months to years, the long-term timing of many pulsars begin to start to show evidence of ‘red’ timing noise. This phenomena manifests as stochastic, low-fluctuation frequency structures in pulsar timing residuals that can be characterised by a red power spectrum. Red noise can arise from various processes that alter the propagation of the radio pulses from the neutron star to our telescopes. These include fluctuations in the density of the interstellar medium along the line of sight ([Keith et al., 2013](#)), and the passage of nanohertz frequency gravitational waves ([Detweiler, 1979](#); [Hellings & Downs, 1983](#)). However, the dominant source of intrinsic red noise over long timescales is believed to be from irregularities in the pulsar rotation period, sometimes referred to as ‘spin noise’.

While the precise mechanism behind pulsar spin noise is unknown, potential external origins include fluctuations in the spin-down torque (Cheng, 1987b), free precession (Stairs et al., 2000; Brook et al., 2014; Kerr et al., 2016), interactions between pulsar magnetospheres and debris (Cordes & Shannon, 2008), undetected planetary companions (Kerr et al., 2015), changes in pulse shape (Brook et al., 2016), and discrete quasi-periodic magnetospheric state switching (Kramer et al., 2006a; Lyne et al., 2010). It can also arise from a variety of internal mechanisms such as the outward exchange of angular momentum due to coupling between the outer crust and superfluid interior (Jones, 1990), and turbulence within the core superfluid (Greenstein, 1970; Link, 2012; Cordes & Shannon, 2008). Long-term studies of large samples of pulsars by Hobbs et al. (2005, 2010) found red noise is common across the pulsar population. They claimed pure random walks in pulsar phase, spin frequency or spin-down rate cannot accurately model the observed structures in the timing residuals (Cordes & Downs, 1985). However, Shannon & Cordes (2010) showed that if the random walk step-sizes are drawn from a power-law function (rather than a Gaussian), then most red noise structures can be replicated.

1.4.3 Glitches

Another type of rotational irregularity that strongly impacts long-term timing are glitches, discrete spin-up events that sometimes decay exponentially over a period of days to months. Glitches are posited to originate from a build-up and release of stress within neutron stars. Potential trigger mechanisms for sudden stress release include starquakes (Ruderman, 1969; Baym et al., 1969) and avalanches of superfluid vortices unpinning from the crustal lattice (Anderson & Itoh, 1975; Alpar et al., 1985; Melatos et al., 2008). A comprehensive review of glitch trigger mechanisms can be found in Haskell & Melatos (2015). There is a distinct bi-modality in the measured glitch sizes (Espinoza et al., 2011c; Yu et al., 2013) with peaks in the fractional step-change in spin-frequency at $\Delta\nu \sim 10^{-8.4}$ Hz and $\Delta\nu \sim 10^{-4.4}$ Hz (Ashton et al., 2017). However, it is unclear whether the observed glitch distribution is complete, as glitches with smaller sizes can be difficult to distinguish from the effects of red noise. While glitches are distinct from red noise, there are suggestions that small ‘micro-glitches’ (Cheng, 1987a) and the long-term re-pinning of vortices following a glitch (Melatos et al., 2008) can contribute to the overall red noise budget of some pulsars.

1.5 The Double Pulsar PSR J0737–3039A/B

The double neutron star system PSR J0737–3039A/B provides one of the best laboratories for performing precision tests of Einstein’s theory of General Relativity and constraining alternate theories of gravity. It was initially discovered in 2003 after the detection of the 22 ms partially-recycled pulsar, PSR J0737–3039A (pulsar A hereafter), by [Burgay et al. \(2003\)](#). Soon after the initial discovery [Lyne et al. \(2004\)](#) reported the detection of radio pulses with a periodicity of 2.8 s originating from the companion neutron star, PSR J0737–3039B (pulsar B hereafter). This is the only known double neutron star system in which both neutron stars had been detected as pulsars, earning it the nickname ‘The Double Pulsar’. Precision timing of both pulsars allowed for independent measurements of the projected semi-major axis of the system to be measured for the first time in any double neutron star. High-precision timing in the years following its discovery enabled the detection of five post-Keplerian parameters and the pulsar mass ratio, which in turn enabled four independent tests of GR to be performed assuming the pulsar masses of $m_A = 1.3381 \pm 0.0007 M_\odot$ and $m_B = 1.2489 \pm 0.0007 M_\odot$ from the overlap between $\dot{\omega}$ and R . All four tests of gravity were found to be consistent with the expected values from GR to a precision of 0.1-0.05% ([Kramer et al., 2006b](#)).

1.5.1 Magnetospheric eclipses

In addition to reporting the discovery of pulsar B, [Lyne et al. \(2004\)](#) noted the radio emission from pulsar A would disappear for 20 to 30 s around superior conjunction (i.e, when pulsar B is in line with pulsar A from our perspective). Observations performed by [Kaspi et al. \(2004\)](#) using the 100-m *Robert C. Byrd Green Bank Telescope* (GBT) revealed pulses from pulsar A were occasionally detected during some eclipses. [McLaughlin et al. \(2004\)](#) re-analysed these GBT observations at a higher time resolution, revealing dramatic variations in the light curve of pulsar A during eclipse, with peaks and troughs that are modulated at once and twice the rotation period of pulsar B. Theoretical modelling of the eclipse light curves by [Lytikov & Thompson \(2005\)](#) demonstrated the intensity variations of pulsar A during the eclipses could be completely described by a simple geometric model. They suggest the radio pulses from pulsar A undergo synchrotron absorption as they propagate through a relativistic pair-plasma trapped within the closed magnetic field lines of pulsar B. In addition to using this model to place the first geometric constraints on the spin and magnetic axes of pulsar B, they also demonstrated the impact geodetic precession ($\Omega_{\text{SO}}^{\text{B}}$; a phenomena induced by relativistic spin-orbit coupling) of pulsar B would

have on the observed light curve. The use of the [Lyutikov & Thompson \(2005\)](#) model for measuring $\Omega_{\text{SO}}^{\text{B}}$ was fully realised by [Breton et al. \(2008\)](#). Using a Bayesian framework to perform joint fits to 63 separate eclipses observed by the GBT over 3.8 years, they recovered a geodetic precession rate of $\Omega_{\text{SO}}^{\text{B}} = (4.77_{-0.65}^{+0.66})^{\circ} \text{ yr}^{-1}$, which agrees with the expected value of $5.0734 \pm 0.0007^{\circ} \text{ yr}^{-1}$ from GR to an uncertainty of 13%. This measurement also enabled a strong-field test of relativistic spin-orbit coupling, in which the observed and GR-predicted values were found to be consistent to within $\sim 10\%$.

As the spin axis of pulsar B continued to precess, the radio beam gradually moved beyond our line of sight, and completely disappeared from view in 2008 ([Perera et al., 2010](#)). Radio pulses originating from pulsar B are not predicted to be detectable until sometime in 2024 ([Breton, 2009](#)). This disappearance meant that continued measurements of the geodetic precession of pulsar B via the joint eclipse fitting method – which relied on an accurate determination of its rotation phase via timing measurements – was no longer feasible. In principle, observations of the eclipses with greater sensitivity than afforded by the GBT could be used to constrain the geometry of pulsar B without the need for a phase-connected solution ([Breton et al., 2008](#)). Such observations are now regularly obtained with the recently commissioned MeerKAT telescope (see figure 9 of [Kramer et al. 2021](#)), the analysis and interpretation of which are the focus of Chapter 7.

1.6 Thesis outline

This thesis focuses on using both new and upgraded radio facilities, as well as programmes with significant legacy datasets, to explore the magnetospheric and rotational properties of pulsars, thereby improving our understanding of these extreme objects. Regular pulsar timing programmes that survey large samples of pulsars, such as those undertaken at the Molonglo Observatory Synthesis Telescope and Parkes radio telescope, are necessary for testing theories on how pulsars should evolve over time. Additionally, the advent of new receiver technologies, such as the Ultra-Wideband Low receiver system on Parkes, and the commissioning of new, extremely sensitive facilities such as MeerKAT, both offer an unprecedented view of the magneto-ionic properties of enigmatic neutron stars such as radio-loud magnetars and the Double Pulsar.

The remaining chapters of this thesis are outlined as follows.

Chapter 2 presents a brief rundown of how pulsars are observed using radio telescopes. This includes an overview of the signal path of a radio telescope, technical details on the three telescopes used for this thesis, and a description of the pulsar timing technique.

Chapter 3 was published as “The UTMOST timing programme II. – Timing noise

across the pulsar population” in *Monthly Notices of the Royal Astronomical Society*. It presents the measurement and characterisation of glitches and pulsar timing noise in 300 pulsars, and the use of hierarchical Bayesian inference to assess how the strength of timing noise depends on pulsar spin frequency and spin-down rate.

Chapter 4 has been published as “The impact of glitches on young pulsar rotational evolution” in *Monthly Notices of the Royal Astronomical Society*. In this work, we cover the impact glitches have had on the rotational evolution of 74 pulsars timed by the Parkes radio telescope over ~ 5 -30 years. We also demonstrate that both glitches and non-glitched pulsars with large apparent braking torques are dominated by the same inter-glitch recovery process. Also included is a noise-limited catalogue of all pulsar glitches (124 in total) within the data set.

Chapter 5 is published as “Spectropolarimetric properties of Swift J1818.0–1607: A 1.4 s radio magnetar” in *The Astrophysical Journal Letters*. We present an analysis of the wideband spectral, polarisation and single-pulse properties of the newly discovered magnetar via Bayesian parameter estimation. Discussed are potential links to more ordinary pulsars, and implications for the origin of this unusual radio-loud magnetar.

Chapter 6 was published as “The dynamic magnetosphere of Swift J1818.0–1607” in *Monthly Notices of the Royal Astronomical Society*. Here we extend the spectral and polarimetric analysis of Chapter 5 to 6-months of Swift J1818.0–1607 monitoring. This includes the detection of a new type of emission mode switching and variable radio emission height in this magnetar. Also discussed is the possible discovery that the radio beam of this magnetar can switch between two discrete, co-located magnetic poles connected via a coronal loop.

Chapter 7 describes our modelling of the Double Pulsar eclipse light curves detected with MeerKAT, along with an investigation of the polarisation properties of pulsar A at different eclipse phases. This includes an updated measurement of the geodetic precession rate of pulsar B and the discovery of birefringence in its magnetospheric plasma.

Chapter 8 summarises the major results of the thesis. Potential future directions are discussed.

2

Observing pulsars with radio telescopes

The focus of this thesis is the analysis and interpretation of pulsar data collected by three large radio observatories. Specifically, the upgraded Molonglo Observatory Synthesis Telescope (UTMOST; [Bailes et al., 2017](#)) and CSIRO 64-m Parkes radio telescope (also known as *Murriyang*), both located in New South Wales, Australia, and the 64 antenna Meer [More] Karoo Array Telescope (MeerKAT; [Jonas, 2009](#)) located in the Northern Cape province of South Africa. This chapter overviews the basic principles behind observing radio pulsars using these facilities.

2.1 Basic signal path

Most pulsars are discovered and observed using telescopes operating at radio frequencies. In this section we describe the most common components of a single-dish radio telescope, including the basic signal path followed, and how the incident radio pulses from a pulsar are recorded and converted to formats used in pulsar timing and searching. Note that much of this section follows the standard pulsar textbook ([Lorimer & Kramer, 2012](#)) hence references are not listed individually unless required. Additionally, this discussion assumes a telescope operating at millimetre to decimetre wavelengths. Telescopes that operate at longer wavelengths (metres to decametres) often have different antenna topology and make use of other data recording approaches than what we describe below (direct sampling versus down-conversion).

A basic radio telescope is made up of two sub-systems depicted in [Figure 2.1](#): the reflector, feed assembly and superheterodyne receiver (the ‘front-end’) and signal digitiser/channeliser (the ‘backend’). The signal path begins with the incident radiation being collected and focussed by the parabolic primary reflector to the prime focus where a receiver assembly may be mounted. Note that positioning the receiver system at the prime

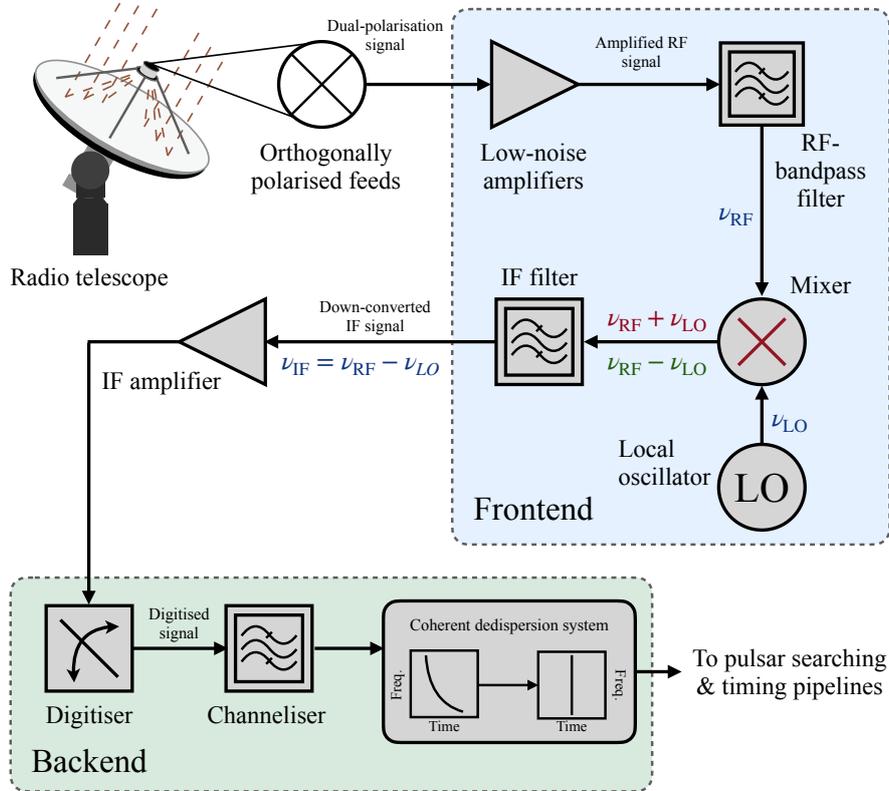


Figure 2.1 Block diagram showing the basic signal path of a standard radio receiver used for observing pulsars. Note that all of the highlighted stages apply to both of the orthogonally-polarised signals output by the receiver.

focus does result in some of the underlying antenna surface being blocked by the feed housing. Many radio telescope designs avoid this issue by using a secondary reflector to send the radio waves to a receiver assembly located at the Cassegrain focus, or by making use of an off-axis Gregorian reflector and feed assembly to ensure the primary reflector is unobstructed. Depending on what frequency range the telescope is intended to operate at, the reflector surfaces are either made of metallic mesh or solid metal panels. The sensitivity of a telescope ('gain') is directly proportional the surface area of its parabolic primary reflector. Telescope gain is defined as the increase in signal strength obtained by focussing the incident radiation from one direction on the sky versus an equally sensitive, isotropic ('all-sky') antenna. The apparent flux of an astronomical object will vary depending on how close it is positioned to the central axis ('boresight') of the primary reflector. This arises from the gain of a radio telescope varying as a function of orientation and distance

from the boresight of the primary reflector, often referred to as the telescope ‘beam shape’. The far-field radiation pattern of a radio telescope is given by the Fourier transform of its aperture, which can be approximated to a 2-dimensional Bessel function (Balanis, 2005). The central maxima (‘lobe’) of the Bessel function forms the ‘primary beam’ of the telescope, the diameter of which is defined by the half-power point or full-width half maximum (FWHM) as

$$\text{FWHM} = 1.22 \frac{\lambda}{D} \text{ rad}, \quad (2.1)$$

where λ is the wavelength of the incident radiation and D is the diameter of the parabolic primary reflector. Other lobes of the Bessel function that appear further out from the primary beam are referred to as ‘side-lobes’.

After being focussed by the reflector(s), the electromagnetic radiation reaches the antenna feed, which typically consists of two orthogonally aligned antennas that are sensitive to either circularly or linearly polarised radio waves. Most receivers use a waveguide to channel the focussed radiation onto the feed, however several modern receiver designs eschew the standard feed-horn in favour of more novel designs that maximise detector bandwidth (e.g. the Parkes Ultra-Wideband Low receiver; Dunning et al. 2015) or field of view on the sky (e.g. phased-array feeds; Hampson et al. 2012). The antenna feed converts the incident radiation into a time-varying electrical signal that is then amplified via a low-noise amplifier (LNA). Both the feed and LNAs maximise the amount of signal recovered at the telescope by minimising the amount of thermal noise in the system (T_{sys}) through careful design and, on occasion, active cooling with cryogenic refrigeration systems. After amplification, the signal is then fed through a bandpass filter to extract the frequency band of interest, before then being passed to the down-conversion system.

Down conversion consists of passing the amplified signal through a heterodyne mixer, where the signal with some input frequency (ν_{RF}) is mixed with a stable, band-limited signal produced by a synchronised local oscillator (ν_{LO}) that is usually driven by the observatory station clock. This produces the heterodyne products, $\nu_{\text{RF}} + \nu_{\text{LO}}$ and $\nu_{\text{RF}} - \nu_{\text{LO}}$ that are then passed through an intermediate-frequency (IF) filter that allows one sideband to be further sent down the signal chain while rejecting other inputs. The resulting IF-filtered signal may undergo another amplification stage at an ‘IF-amplifier’. From here, the signal is sent to the telescope backend where it passes through an analogue-to-digital converter (the ‘digitiser’) and is ‘real-sampled’ at the Nyquist frequency ($\nu_{\text{Nyq}} = 2\nu_{\text{IF}} = 1/2\Delta\nu$, where $\Delta\nu$ is the bandwidth of the signal) in the case of a single digitiser, or ‘complex sampled’ at half ν_{Nyq} when a pair of digitisers are used. The sampled ‘voltages’ are then passed to the channeliser, which in many modern radio telescopes consists of a

digital polyphase filterbank that often use field programmable gate arrays to perform a weighted Fourier transform of the signal to produce a specified number of voltage timeseries for each frequency channel (a ‘voltage filterbank’). These channelised data are typically sent to a computing server where it is subsequently converted into either pulsar timing or search-mode data products.

Pulsar search-mode data is produced by passing the voltages through a square-law detector (usually implemented in software) to form an array of bandpass filters that each carry a single frequency sub-band (termed a ‘filterbank’). Filterbanks used in pulsar astronomy usually follow either the SIGPROC¹ (Lorimer, 2011) or PSRFITS (Hotan et al., 2004) formats. From here, it can then be passed to pulsar (or FRB) search algorithms that look for periodic signals (or bursts) at various dispersion trials (e.g. PRESTO; Ransom 2001). For pulsar-timing (‘fold-mode’) observations, the voltage data are passed to a dedispersion algorithm, such as DSPSR² (van Straten & Bailes, 2011), that removes the frequency-dependent dispersive delay introduced to the signal by the ISM (see Section 1.2.1 of Chapter 1) via coherent dedispersion at the known DM of the pulsar (Hankins, 1971; Hankins & Rickett, 1975). The resulting dedispersed filterbank is then folded at the predicted topocentric period of the pulsar using a supplied ephemeris, and saved to a PSRCHIVE-format³ archive (Hotan et al., 2004; van Straten et al., 2012). These folded archive files form the basic data products that are used in most pulsar studies.

The sensitivity of a telescope to radio pulses from a pulsar can be evaluated via the radiometer equation (Lorimer & Kramer, 2012), where the resulting signal-to-noise ratio (S/N) for a source with period P and pulse-width W is given by

$$S/N = \frac{GS_{\text{mean}}\sqrt{N_p\Delta\nu t_{\text{obs}}}}{T_{\text{sys}}}\sqrt{\frac{P-W}{W}}. \quad (2.2)$$

Here, $G = A_{\text{eff}}/2k_{\text{B}}$ is the telescope gain (in K Jy^{-1}), which depends on the Boltzmann constant (k_{B}) and the effective telescope aperture (A_{eff}), while S_{mean} is the average flux density of the source in Jy , N_p is the number of detected orthogonal polarisations (either 1 or 2), $\Delta\nu$ is the telescope bandwidth in Hz , t_{obs} is the observing timespan in seconds and T_{sys} is the total system temperature in K . While this equation provides a useful metric for predicting the general performance of a telescope, it does leave out a number of important factors that must also be taken into consideration when planning pulsar surveys and pulsar monitoring/timing programmes, or designing new pulsar facilities. This includes the field

¹<http://sigproc.sourceforge.net/>

²<http://dspsr.sourceforge.net/>

³<http://psrchive.sourceforge.net/>



Figure 2.2 The CSIRO Parkes 64-m radio telescope. Image taken by Marcus Lower.

of view of the telescope, which as noted earlier in equation 2.1 is strongly dependent on both the primary reflector size (assuming a single parabolic dish) and observing frequency, and the telescope response time, i.e how quickly it can move from one source to another.

2.2 The Parkes radio telescope: A single-dish instrument

The CSIRO Parkes radio telescope, pictured in Figure 2.2, is a large, 64-m single dish antenna located north of the town of Parkes in New South Wales, Australia. It consists of the aforementioned 64-m reflector attached to an altitude-azimuth drive situated on a turret-style tower. The focus-cabin positioned at the prime focus houses up to three receiver assemblies on a movable track, allowing them to be swapped between subsequent observations. There is also a range of other receiver systems that are stored on-site, which can be swapped with a receiver in the focus cabin during scheduled maintenance periods.

Parkes has long and storied history of being at the forefront of pulsar research. Well over 1200 individual pulsars have been discovered in various surveys undertaken at Parkes over the past 50-years (Komesaroff et al., 1973; Manchester et al., 1978; Johnston et al., 1992a; Manchester et al., 1996; Lyne et al., 1998; Burgay et al., 2006; Edwards et al.,

2001; Jacoby et al., 2009; Keith et al., 2010; Keane et al., 2018), 781 of which were detected during the highly successful Parkes Multibeam Pulsar Survey alone (Manchester et al., 2001; Lorimer et al., 2006). Alongside extensive pulsar surveys, Parkes has also contributed a number of firsts to the field of pulsar astronomy. This includes the first detection of a pulsar glitch (Radhakrishnan et al., 1969), the first pulsar found orbiting a main-sequence star (Johnston et al., 1992b), the first double pulsar (Lyne et al., 2004), and discoveries of entirely new classes of radio pulsars in the form of radio-loud magnetars (Camilo et al., 2006) and rotating radio transients (RRATs; McLaughlin et al. 2006). The telescope is also currently host to a set of large-scale pulsar timing experiments. The Parkes Pulsar Timing Array (PPTA; Manchester et al. 2013) aims to use the high-precision timing of an ensemble of millisecond pulsars to detect nanohertz-frequency gravitational waves originating from inspiralling supermassive black hole binaries (e.g. Goncharov et al., 2021b), and (potentially) other exotic gravitational-wave sources (e.g. cosmic strings; Yonemaru et al. 2021). Chapter 4 of this thesis uses data collected by the Parkes young-pulsar timing programme, which was initially designed to support the detection of gamma-ray emitting pulsars with the *Fermi* satellite (Smith et al., 2008; Weltevrede et al., 2010; Johnston et al., 2021), while Chapter 6 analysed data that was taken under the Parkes magnetar monitoring programme.

Starting in the early 2000s, pulsar studies at Parkes were largely performed at wavelengths of 20 cm with the primary beam of the Parkes multi-beam and H-OH receivers (Staveley-Smith et al., 1996; Granet et al., 2011), and at both 10- and 50-cm by the dual-band ‘10/50-cm’ receiver (Granet et al., 2005). Note, the 10/50-cm receiver was re-tuned in 2009 to cover the 40-cm band due to increased levels of RFI within the original 50-cm band. Signals produced by these receivers were down converted to IF within the focus cabin, before being band-pass filtered and amplified in the receiver control room. From here they were sent to a number of signal processing backends that were developed and operated at Parkes over several decades (see section 2 of Hobbs et al. 2011 for an overview of most ‘legacy’ pulsar backends). The dataset used in Chapter 4 was primarily captured using pulsar observations taken by the Parkes multi-beam and H-OH receivers that were subsequently digitised by the Parkes Digital Filter Bank (PDFB) series of backends (Manchester et al., 2013).

In 2018, a new wideband receiver system and backend was installed at Parkes. The Ultra-Wideband Low (UWL) receiver operates over a continuous frequency range from 704-4032 MHz, effectively replacing the multi-beam, H-OH and 10/50-cm receivers with a single observing system (Hobbs et al., 2020). It consists of a cryogenically cooled feed,

LNAs and RF-amplifiers, with digitisation of the RF signal taking place in the focus cabin, before being streamed to the MEDUSA backend in the telescope tower for further processing. Pulsar fold-mode and search-mode data can also be captured from the UWL via the PDFB4 backend over a maximum frequency range of 1024 MHz. Data analysed in Chapters 5 and 6 were collected by the UWL receiver and MEDUSA backend.

2.3 UTMOST and MeerKAT: Interferometers

The two other facilities used to capture data for this thesis are the Molonglo Observatory Synthesis Telescope (MOST) and MeerKAT. Unlike Parkes, these two telescopes consist of arrays of smaller telescopes that use the principles of interferometry to coherently add their received signals to form effective single instruments with large collecting areas and resolving power.

UTMOST

The Molonglo Observatory Synthesis Telescope, a portion of which is depicted in Figure 2.3, is located in New South Wales, Australia, approximately 40 km east of Canberra. It consists of two 1.6 km long half-cylindrical reflectors that aligned are East-West and North-South in a Mills-cross design. This telescope has a long history of being used for pulsar research. Starting in its initial design as the 408 MHz Molonglo Cross Telescope, the first Molonglo pulsar surveys discovered 186 pulsars (Large et al., 1968b; Large & Vaughan, 1971; Manchester et al., 1978), including the Vela pulsar, PSR J0835–4510 (Large et al., 1968a). Additional pulsar surveys (Manchester, 1984; D’Amico et al., 1985, 1988) and a timing programme (Siegman et al., 1993) were conducted with the telescope following the conversion to its current 843 MHz operating frequency. This was in parallel to its primary use as a synthesis imaging array, which produced many of the premier surveys and catalogues of the Southern radio sky (Whiteoak & Green, 1996; Bock et al., 1999; Green et al., 1999; Mauch et al., 2003; Murphy et al., 2007; Green et al., 2014). Starting in 2013, the MOST underwent a major upgrade under the UTMOST project, which transformed the telescope into a pulsar-timing and radio transient detecting machine (Bailes et al., 2017).

At the time of writing, both the East-West and [a portion of the] North-South arms are operational as semi-independent arrays as part of the UTMOST-2D project. The UTMOST data that was used in this thesis was collected using only the East-West arm of the telescope, hence the remainder of this section focuses on describing this portion of



Figure 2.3 The East-arm of the Molonglo Observatory Synthesis Telescope. Image taken by Marcus Lower.

the observatory. The East-West arm consists of 352 individual subsections (colloquially referred to as ‘modules’), which when combined form a large aperture synthesis array with an 18000 m^2 collecting area. Each module contains 7744 circular dipoles (‘ring antennas’) that operate at a central frequency of 835 MHz with 30 MHz of bandwidth. These antennas are only sensitive to right-hand circularly polarised radiation. Originally, the telescope could be pointed in both the north-south and east-west directions by mechanically rotating the cylinder about its long axis in the north-south direction and rotating the feed elements to steer the phased beam east-west. However, in mid-2017 the telescope was converted to a meridian transit instrument after the mechanism responsible for rotating the ring antennas was decommissioned in the interests of alleviating strain on the system. Data from each module are coherently summed in phase to form up to four tied-array beams that can be used to track pulsars as they transit the primary beam of the telescope, which covers an area of $4.25^\circ \times 2.8^\circ$ on the sky. Specific details on the general observing system can be found in [Bailes et al. \(2017\)](#), with additional information covering the pulsar timing programme in [Jankowski et al. \(2019\)](#) and Chapter 3 of this thesis.



Figure 2.4 A handful of the 64 antennas located in the MeerKAT core. Image credit: South African Radio Astronomy Observatory.

MeerKAT

MeerKAT is the Square Kilometer Array (SKA)-mid precursor telescope located in the Karoo Desert of South Africa’s Northern Cape province. It consists of 64 individual telescopes (a handful of which are shown in Figure 2.4), each of which is comprised of a 13.5 m dish with an offset arm containing the secondary reflector and receiver system in a ‘feed-down’ Gregorian optical configuration (Jonas & the MeerKAT Team, 2016). The array has two fully operational, cryogenically-cooled receiver fleets: the ‘Ultra-High Frequency (UHF)’ and ‘L-band’ systems that cover the 580-1080 MHz and 900-1650 MHz bands respectively. A third set of ‘S-band’ receivers supplied by the Max Planck Institute für Radioastronomie (Kramer et al., 2016) that cover the 1750-3000 MHz band are currently undergoing commissioning. Over half of the antennas within the array are arranged into a dense core, where 44 out of 64 antennas are located within a 1 km diameter circle, while the remaining 20 antennas are distributed out to a distance of 4 km from the array centre for a maximum baseline of 8 km. Signals captured by the receivers are digitised at the antennas, skipping the superheterodyne mixing step in Figure 2.1. The signals are then sent along buried optical fibres to the Karoo Array Processing Building, located ~ 10 km from the telescope site, where they are processed by the correlator/beamformer digital signal processor. Following correlation and beamforming, the resulting correlated voltage

time series is sent to the Pulsar Timing User Supplied Equipment (PTUSE) backend, a prototype of the pulsar signal processor that will ultimately be deployed on the SKA. PTUSE consists of four individual servers, each capable of capturing simultaneous pulsar search- and fold-mode data, which when combined with the ability to split MeerKAT into subarrays of varying antenna configurations, enables a wide variety of possible observing modes. Additionally, high-level technical details of both the telescope system, PTUSE backend and data calibration procedures are described in [Bailes et al. \(2020\)](#).

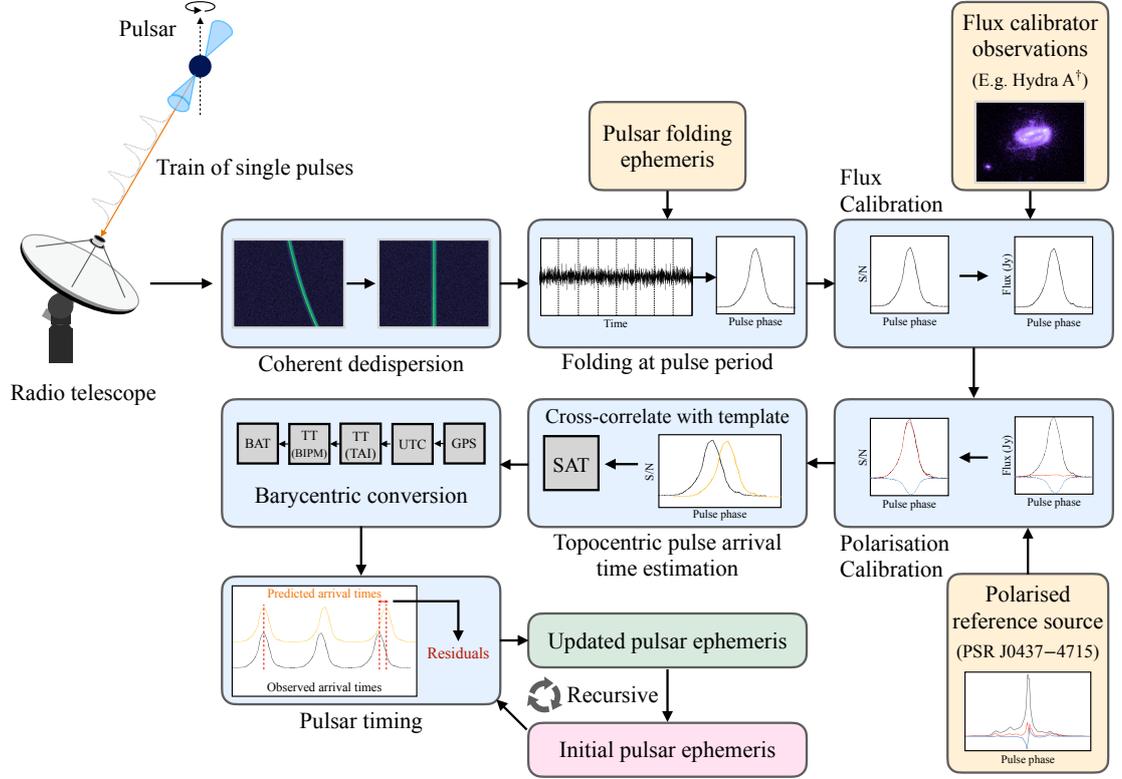
2.4 Pulsar timing

The most basic properties that can be determined from observations of pulsars are their spin frequency, ν (or rotation period, P), and spin-down rate ($\dot{\nu}$ or \dot{P}). As discussed in Chapter 1, we can use these two properties to obtain estimates of a pulsar’s spin-down luminosity, age and surface magnetic field strength. While the tracking and fitting of ν over time can enable $\dot{\nu}$ to be measured, most pulsar properties are typically inferred via pulsar timing. The basic steps to time a pulsar are shown in Figure 2.5, which we expand upon below.

The primary goal of pulsar timing is to keep a precise record of every rotation a pulsar has undergone from the beginning to the end of a timing dataset. This is done by comparing the measured pulse time of arrival (ToA) with the predicted arrival time from a timing model, and then performing some form of regression to update the timing model with improved measurements of pulsar properties, or in some cases the addition of newly detected parameters. ToAs are typically computed by cross-correlating the time-tagged pulsar data with a standard profile, either a smoothed high S/N observation obtained after many hours of integration or an analytic template generated from a fit to the observed profile. This technique, known as template matching, generally assumes the observed pulse profile is a scaled version of the template, $S(t)$, shifted by a small amount of time, τ , and superimposed on radiometer noise, $N(t)$, as

$$O(t) = AS(t - \tau) + B + N(t) \quad (2.3)$$

where A and B are constants. A number of methods for performing this template matching technique have been devised, including the Fourier Phase Gradient ([Taylor, 1992](#)), Gaussian Interpolation Shift ([Hotan et al., 2005](#)) and Fourier-domain Monte-Carlo (FDM) methods. The FDM method as implemented in the PSRCHIVE ([Hotan et al., 2004](#); [van Straten et al., 2012](#)) pulsar software library under the PAT tool was used extensively



[†]Hydra A image: NASA/ESA, G. Tremblay, R. Mittal

Figure 2.5 Diagram of the standard processes behind observing, processing and timing a pulsar.

throughout this thesis. The measured time offset τ is the pulse arrival time at a telescope on Earth, t_{topo} (‘topocentric’ site arrival time; SAT), which can be related to an arrival time at the Solar-System Barycenter (SSB), t_{SSB} (barycentric arrival time; BAT), as (Hobbs et al., 2006)

$$t_{\text{SSB}} = t_{\text{topo}} + \Delta t_{\text{clk}} - \Delta t_{\text{DM}} + \Delta t_{\text{R}_{\odot}} + \Delta t_{\text{S}_{\odot}} + \Delta t_{\text{E}_{\odot}}. \quad (2.4)$$

The first correction (t_{clk}) incorporates corrections to the observatory clock and the necessary transformations to synchronise it to the Terrestrial Time (TT(TAI)) standard (typically preceded by a conversion to GPS time, then Universal Coordinated Time; UTC), which maintained by the Bureau International de Poids et Mesures (BIPM) as the average time from an ensemble of atomic clocks distributed around the world. The second correction is the removal of the frequency-dependent delay caused by dispersion (Δt_{DM} ; see Equation 1.13 in the previous chapter). Third, is the Römer delay ($\Delta t_{\text{R}_{\odot}}$), which is the classical light-travel time from the Earth to the SSB. This particular correction takes into

account a conversion from the telescope position on the Earth’s surface to the geocenter, the precise motion of the Earth around the sun and the positions of all minor bodies in the Solar System. The final two terms are the general relativistic corrections to the light-travel time induced by massive bodies within the Solar System. These include the Shapiro delay (Δt_{S_\odot} ; [Shapiro 1964](#)), originating from the increased effective distance the radio pulses have travelled due to the curvature of spacetime around the Sun and planets, and the Einstein delay (Δt_{E_\odot}) resulting from the combined effects of time dilation and gravitational redshift from the potential wells of the Sun, planets and the Earth. From here, the BATs are then subtracted from the predicted pulse arrival times generated from a timing model to obtain the ‘timing residuals’ of the pulsar.

Structures in the timing residuals arise from deviations from the expected pulse arrival times, which can result from either incorrectly modelled pulsar parameters or new unmodelled physics. For instance, the delays in Equation 2.4 are position dependent, hence an incorrect sky-position will result in a characteristic sinusoidal structure in the timing residuals with a periodicity of one year. This can also be exploited to infer the distances to relatively nearby pulsars by using the long-baseline of the Earth’s orbit around the Sun to measure their parallaxes, which result in a 6-month periodicity in the residuals. Pulsars with large transverse velocities will exhibit an annual sinusoid with a growing amplitude due to the changing position induced by their proper-motion across the sky. The presence of binary companions orbiting a pulsar will also introduce additional structures in the timing residuals, which are expanded upon further below. Changes in the pulsar rotation period due to sudden timing events, such as glitches, or stochastic wandering from timing noise lead to the pulses arriving earlier (or later) than expected. This can require re-fitting of the ToAs with more complex timing models via the use of robust Bayesian inference techniques that simultaneously model the astrometric, rotational and stochastic timing properties of pulsars. The use of these techniques in understanding rotational irregularities in pulsars are explored in Chapters 3 and 4.

2.4.1 Binary pulsars

Pulsars that reside in binaries or are host to a planetary system (e.g. [Wolszczan & Frail, 1992](#)) require the addition of Keplerian orbital parameters to the timing model in order to account for the reflex motion of the pulsar around its companion(s). The extra Keplerian parameters are accounted for by adding four additional time-correction terms to Equation

2.4, the sum of which is given by (Damour & Taylor, 1992)

$$\Delta t_{\text{BIN}} = \Delta t_{\text{RB}} + \Delta t_{\text{SB}} + \Delta t_{\text{EB}} + \Delta t_{\text{AB}}. \quad (2.5)$$

Here, Δt_{RB} is the binary Römer delay imparted by the Keplerian orbital motion of the system, Δt_{SB} is the binary Shapiro delay caused by the passage of the pulsar radiation through the curved spacetime around the companion at superior conjunction, Δt_{EB} is the binary Einstein delay from the gravitational time-dilation and redshift from the companion, and Δt_{AB} accounts for the special relativistic aberration induced by binary motion.

These delays can be geometrically related to a set of physical Keplerian and relativistic (post-Keplerian; PK) parameters. The Keplerian parameters are listed below as:

- The orbital period of the binary (P_b),
- orbital eccentricity (e),
- projected semi-major axis of the pulsar, $x = a \sin i/c$, in which a is the semi-major axis of the orbit, i the binary inclination angle and c the vacuum speed of light,
- longitude of the ascending node (Ω_{asc}),
- longitude of periastron (ω_0),
- and finally, the time of periastron passage (T_0).

For pulsars residing in non-relativistic binaries, these parameters are often sufficient enough to completely describe their orbital configurations. A schematic diagram showing how these angles relate to the geometric orientation of a binary system is shown in Figure 2.6. However, for pulsars that do reside in relativistic systems, a set of PK parameters that approximate the effects of general relativity are often required. These PK-parameters include:

- precession of and nutation of the longitude of periastron (k), the lowest post-Newtonian contribution to which is periastron advance ($\dot{\omega}$),
- special and general relativistic time-dilation and gravitational redshift (γ),
- orbital period decay due to gravitational-wave radiation (\dot{P}_b),
- Shapiro-delay amplitude or range (r) and shape ($s \equiv \sin i$), the latter of which indicates the extent of the Shapiro delay in orbital phase,

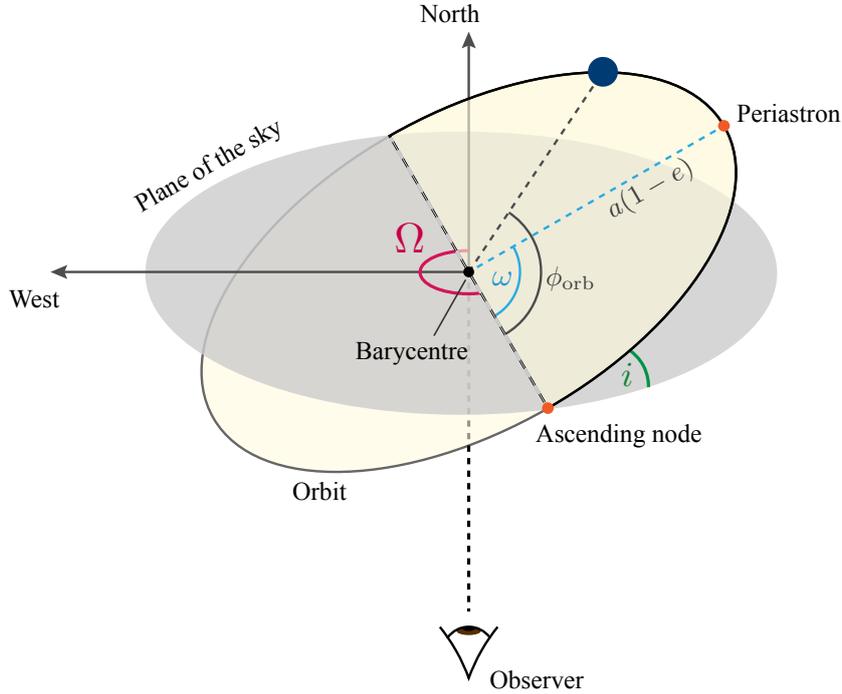


Figure 2.6 Schematic view of a Keplerian orbit. The intersection of the orbital (yellow) and sky (grey) planes is the line of nodes (black-dashed line), where the orientation of the orbit as seen by the observer is given by the longitude of the ascending node in the plane of the sky (Ω). It is tilted from the plane of the sky by an inclination angle i . Periastron marks the point of closest approach between the pulsar and its companion at a distance of $a(1 - e)$, for semi-major axis a and eccentricity e . The longitude of periastron (ω) is defined as the angle separating periastron from the ascending node, while the orbital phase (ϕ_{orb}) is the angle between the ascending node and the position of the pulsar.

- relativistic deformation of the orbit (parameterised by δ_r and δ_θ),
- relativistic aberration and retardation (A and B),
- and orbital eccentricity decay (\dot{e}).

There is also the secular evolution of the projected semi-major axis (\dot{x}), which can arise from both relativistic (e.g. Lense-Thirring precession [Venkatraman Krishnan et al. 2020](#)) and non-relativistic changes in the geometry of a binary system. Approximating the pulsar and its companion to point-masses, each of these parameters can be written as functions of the pulsar and companion masses (m_p and m_c), as well as the corresponding Keplerian orbital elements. Some PK parameters either cannot be measured via traditional pulsar timing, such as the relativistic orbital deformation parameter δ_r which is covariant with combinations of orbital and spin parameters ([Damour & Taylor, 1992](#)), or are strongly

covariant with other PK parameters. As an example of the latter, geodetic precession induced by spin-orbit coupling (Ω_{SO} ; [Damour & Ruffini 1974](#)) may induce a measurable effect on the timing of relativistic pulsars through variations in the aberration parameters, A and B . However, it can in some instances be quantified by its far stronger, indirect impact on the observed pulse profile shape and polarisation, as it causes different sections the pulsar emission cone to cross our line of sight. In extreme cases, it may even cause the pulsar radio beam to completely miss the Earth, resulting in the pulsar disappearing from view for decades to centuries, as occurred with PSR J0737–3039B ([Perera et al., 2010](#); [Noutsos et al., 2020](#)). Measurements of two or more PK parameters can be used to place tight constraints on m_p and m_c , which can in-turn be used to check the consistency of Einstein’s general theory of relativity and constrain the equation of state of nuclear matter (see [Damour & Taylor 1992](#) and [Demorest et al. 2010](#) for additional details).

3

Timing noise across the pulsar population

While pulsars possess exceptional rotational stability, large scale timing studies have revealed at least two distinct types of irregularities in their rotation: red timing noise and glitches. Using modern Bayesian techniques, we investigated the timing noise properties of 300 bright southern-sky radio pulsars that have been observed over 1.0-4.8 years by the upgraded Molonglo Observatory Synthesis Telescope (MOST). We reanalysed the spin and spin-down changes associated with nine previously reported pulsar glitches, report the discovery of three new glitches and four unusual glitch-like events in the rotational evolution of PSR J1825–0935. We develop a refined Bayesian framework for determining how red noise strength scales with pulsar spin frequency (ν) and spin-down frequency ($\dot{\nu}$), which we apply to a sample of 280 non-recycled pulsars. With this new method and a simple power-law scaling relation, we show that red noise strength scales across the non-recycled pulsar population as $\nu^a |\dot{\nu}|^b$, where $a = -0.84_{-0.49}^{+0.47}$ and $b = 0.97_{-0.19}^{+0.16}$. This method can be easily adapted to utilise more complex, astrophysically motivated red noise models. Lastly, we highlight our timing of the double neutron star PSR J0737–3039, and the rediscovery of a bright radio pulsar originally found during the first Molonglo pulsar surveys with an incorrectly catalogued position.

This chapter has been published in *Monthly Notices of the Royal Astronomical Society* (Lower et al., 2020b)

3.1 Introduction

The pulsar timing programme of the UTMOST¹ project (Bailes et al., 2017) monitors more than 400 pulsars using the upgraded Molonglo Observatory Synthesis telescope. This programme runs in parallel with searches for undiscovered pulsars and single pulses from

¹Not an acronym.

rotating radio transients (RRATs), and Fast Radio Bursts (FRBs). These searches have already led to the discovery of thirteen FRBs (Caleb et al., 2017; Farah et al., 2018, 2019; Gupta et al., 2019a,c,b,d) and the intermittent pulsar candidate PSR J1659–54 (Venkatesan et al., 2020). As part of the timing programme, updated rotational models for 205 pulsars were published in the first paper of this series (Jankowski et al., 2019) and 9 pulsar glitches have been reported so far (Jankowski et al., 2015b,a, 2016, 2017; Lower et al., 2018, 2019).

While pulsars are renowned for their capacity to be used as astrophysical clocks, many have been observed to exhibit an intrinsic ‘timing noise’ in the measured arrival times of their pulses. Timing noise – stochastic wandering in observed pulse arrival times – manifests as either a ‘white’ noise processes if the power is distributed normally across all fluctuation frequencies, or ‘red’ noise if the timing residuals are dominated by low-fluctuation frequency structures. White timing noise can arise from instrumental artefacts, unaccounted radio frequency interference (RFI) and pulse-to-pulse shape variations, often referred to as pulse jitter (Staelin & Reifenstein, 1968; Jenet et al., 1998). While fluctuations in the density of the interstellar medium along the line of sight (Keith et al., 2013), and the passage of nanohertz frequency gravitational waves (Detweiler, 1979; Hellings & Downs, 1983) manifest themselves as red noise in pulsar timing residuals, the dominant source of intrinsic red noise over long timescales is believed to arise from irregularities in pulsar rotation. One of two primary sources of rotational irregularities are pulsar glitches, discrete spin-up events that often decay exponentially over the following months to years. Glitches are posited to originate from either the release of built up tension within the neutron star’s crust via starquakes (Ruderman, 1969; Baym et al., 1969), or the unpinning of superfluid vortices from the crustal lattice (Anderson & Itoh, 1975; Alpar et al., 1985; Melatos et al., 2008). The other main type of rotational irregularity is ‘spin noise’, long-term variations in pulsar spins characterized by a red power spectrum (hereafter referred to as red noise). While red noise is distinct from glitches, it may be possible that glitch recoveries and switching between emission/spin-down states contributes to the overall red noise seen in some pulsars. The nature of the relationship between glitches and red noise is also open for debate (Hobbs et al., 2010; Parthasarathy et al., 2019).

While the precise mechanism behind pulsar red noise is unknown, potential external causes include fluctuations in the spin-down torque (Cheng, 1987b), free-precession (Stairs et al., 2000; Brook et al., 2014; Kerr et al., 2016), asteroid belts or debris disks interacting with the pulsar magnetic field (Cordes & Shannon, 2008), undetected planetary companions (Kerr et al., 2015), changes in pulse shape (Brook et al., 2016) and discrete quasi-

periodic magnetospheric state switching (Kramer et al., 2006a; Lyne et al., 2010). Internal mechanisms such as the outward exchange of angular momentum due to coupling between the outer crust and superfluid interior (Jones, 1990), undetected micro-glitches (Cheng, 1987a), vortex re-pinning during glitch recovery (Melatos et al., 2008) and superfluid turbulence (Greenstein, 1970; Link, 2012; Cordes & Shannon, 2008) have also been proposed as mechanisms behind red noise. Long-term studies of large samples of pulsars by Hobbs et al. (2005) and Hobbs et al. (2010) found red noise is common across the pulsar population, and claimed pure random walks in pulsar phase, spin frequency or spin-down rate cannot accurately model the observed structures in the timing residuals (Cordes & Downs, 1985). However, Shannon & Cordes (2010) showed that if the random walk step-sizes are drawn from a power-law function (rather than a Gaussian), then most red noise structures can be replicated.

In this paper we undertake a study of the rotational properties of 300 pulsars that have been observed over the past 1.0-4.8 years by UTMOST. This includes a full reanalysis of nine previously reported glitches (Jankowski et al., 2015b,a, 2016, 2017; Espinoza et al., 2011a; Sarkissian et al., 2017; Palfreyman et al., 2018; Lower et al., 2018; Liu et al., 2019; Sarkissian et al., 2019; Kerr, 2019), while accounting for the effects of timing noise and the discovery of two new glitches. To parameterize the effects of red noise on the timing residuals, and to avoid biasing our measurements of pulsar spin and spin-down, we employ the Bayesian pulsar timing software `TEMPONEST` (Lentati et al., 2014). We search for correlations between pulsar properties and red noise strength, in addition to how it scales as a function of spin and spin-down frequencies across the population.

In Section 3.2 we outline the observing and data processing strategies. In Section 3.3 we describe the phenomenology behind characterising pulsar timing noise and the statistical framework we use to parameterise red timing noise and perform simultaneous measurements of pulsar spin properties. We report on our updated timing models and present the results of our red noise search and glitch analysis in Section 3.4. In Section 3.5 we examine potential links between red noise strength and pulsar properties, in addition to outlining and comparing a new, robust method for determining how timing noise scales across the population. Lastly, in Section 3.6 we draw our conclusions and comment on future applications of our Bayesian framework.

3.2 Observations

3.2.1 System overview

The UTMOST project began with the backend upgrade to the refurbished Molonglo Observatory Synthesis Telescope (Bailes et al., 2017). MOST is a Mills-Cross design aperture synthesis telescope situated approximately 35 km South-East of Canberra, Australia. It is comprised of two 778 m long East-West arms that can be slewed in the north-south direction, and a static North-South arm, that is being re-engineered as part of the UTMOST-2D project (Day et al. in prep.). The telescope operates at a central frequency of 835 MHz² covering a bandwidth of 31.25 MHz. The ring-shaped design of the antenna elements means the instrument is mainly sensitive to right-hand circularly polarized emission.

For the first two years of the timing programme we were capable of mechanically tracking sources in hour angle on the sky. However, maintenance issues and an associated degradation of performance ultimately led us to convert the telescope into a meridian transit instrument in June 2017. While we are no longer able to track sources mechanically, we are able to electronically track up to four pulsars simultaneously as they transit the telescope’s $4^\circ \times 2^\circ$ primary beam. A typical timing observation lasts between 5-20 minutes depending on the brightness of the pulsar and whether it displays interesting behaviours (e.g. emission state-switching). Observations are usually performed autonomously via the scheduler developed for the UTMOST multi-epoch Survey of Magnetars, Intermittent pulsars, RRATs and FRBs (SMIRF: Venkatraman Krishnan et al., 2020), which has improved the efficiency of the timing programme since its June 2017 introduction. Manual observations of targets of interest, phase calibrators and long FRB transit searches are usually performed using the AUTOMATIC MODE scheduler detailed in Jankowski et al. (2019).

3.2.2 Radio frequency interference

Observations conducted by UTMOST are often contaminated by radio frequency interference (RFI) as its frequency band is shared by radio transmissions from two Australian mobile telecommunications providers. As the telescope is an array, voltage addition in phase only occurs for radio emitting sources that are more than a Fresnel scale away (~ 10000 km) from the telescope, while anything closer is attenuated. Although this does reduce the overall amount of observed RFI, it is still prevalent in a significant fraction of observations. Removal of RFI is performed by passing the data through an excision

²The sensitivity of the MOST peaks at ~ 843 MHz, as this is where the resonant cavities are tuned to.

pipeline that involves spectral kurtosis prior to folding of the raw data with DSPSR³ (van Straten & Bailes, 2011), followed by median filtering of the folded archives via the tools in PSRCHIVE⁴ (Hotan et al., 2004; van Straten et al., 2012). Manual RFI removal with PSRCHIVE’s interactive PAZI tool is undertaken when necessary. More recently, we have modified the RFI cleaning pipeline to use CLFD⁵ (Morello et al., 2018), which uses Tukey’s rule to find and zero-weight data corresponding to outliers in the standard deviation, peak-to-peak difference and second bin of the Fourier transform of each sub-integration and channel of a folded observation. This alone has improved the timing accuracy of many slow pulsars we observe by a factor of two. The amount of data lost to RFI excision is typically on the order of 5 percent, but can be as high as 10 to 15 percent during times of high road traffic (and hence an increased number of mobile handsets) near the telescope.

3.2.3 Pulsar-timing dataset

We began the pulsar-timing programme during October 2015 after phasing of the telescope became routine. Limited pulsar observations prior to this date were undertaken while the telescope was still undergoing upgrades and commissioning, but are largely of lower quality when compared to more recent data. A general overview of the UTMOST timing programme can be found in Jankowski et al. (2019), which includes the first scientific results of the timing programme: a study of pulsar proper motions, transverse velocities, pulse duty cycles and flux densities at 843 MHz, and updated rotational and astrometric parameters for 205 pulsars. Currently we perform regular radio monitoring and timing of 412 pulsars, each of which was selected from an initial list of every pulsar for which an observation had been attempted by UTMOST. This includes monitoring the pulsed radio emission of two radio loud magnetars, PSR J1622–4950 and XTE J1810–197. Each pulsar observation typically lasts between 5-20 minutes, depending on the source flux density and declination. After RFI excision, the observations are then summed in frequency and time to produce averaged pulse profiles. These are then cross-correlated with a standard profile, a template generated from a smoothed, high signal-to-noise profile obtained after many hours of integration, to measure the pulse time of arrival (ToA) at the telescope (Taylor, 1992). This ‘topocentric’ ToA is then converted to the ToA at the Solar System Barycentre via the Jet Propulsion Laboratory’s DE430 planetary ephemeris (Folkner et al., 2014). Due to sensitivity limitations of the telescope, most of these pulsars are bright, isolated southern pulsars with relatively long rotation periods. Their basic ob-

³dspsr.sourceforge.net

⁴psrchive.sourceforge.net

⁵github.com/v-morello/clfd

servational parameters are drawn from the Australia Telescope National Facility (ATNF) pulsar catalogue (PSRCAT; Manchester et al. (2005))⁶ and Jankowski et al. (2019), where the spin period, position and DM determination epoch is MJD 57600. Fig. 3.1 shows the spin period/period-derivative ($P-\dot{P}$) diagram for the full set of pulsars regularly monitored by UTMOST.

The cadence with which we observe individual pulsars varies from days to months, depending on their physical properties, and whether they exhibit interesting behaviour such as nulling, glitches or mode-changing. Precise observation cadences and lengths are defined by a pulsar’s position in the sky, apparent brightness and the scientific benefit of performing observations with daily to monthly cadences. We provide this information to SMIRF, which autonomously schedules and performs the observations. Limiting the amount of mechanical wear on the telescope due to slewing is also factored into how often particular pulsars are observed.

3.3 Pulsar-timing strategy

To determine the timing properties of a pulsar, we must first establish a phase-connected timing solution. Pulsars that have undergone glitches or exhibit excess structures due to timing noise are difficult to solve coherently over long timing baselines, often requiring the addition of discrete phase jumps before updating the timing model. Once we have a phase coherent solution, we use TEMPO2 (Hobbs et al., 2006) to assign relative pulse numbers to each ToA, which we then track to avoid phase wraps when attempting to update the timing model. The effects of timing noise can be mitigated by including fits for higher order polynomials, corresponding to higher order spin-derivatives, into the timing model, or by subtracting a series of harmonically related sinusoids (e.g. FITWAVES in TEMPO2; Hobbs et al. (2006)). While these methods are useful for pre-whitening the timing residuals, they assume the measured pulsar properties and timing noise are uncorrelated. Attempts to avoid biases induced by correlations in the timing residuals include using a transform of the covariance matrix based on Cholesky decomposition to whiten the timing residuals (Coles et al., 2011), which enables the timing residual power spectrum to be fit by a steep red power-law. However, this method requires some *a priori* knowledge of the covariance matrix, and that the correlated timing noise process is stationary in the post-fit timing residuals. van Haasteren & Levin (2013) showed that the assumption of stationarity breaks down during the fitting process, leading to incorrect uncertainties on the spectrum estimates, and an incorrect covariance matrix for the pulsar timing model.

⁶www.atnf.csiro.au/research/pulsar/psrcat/

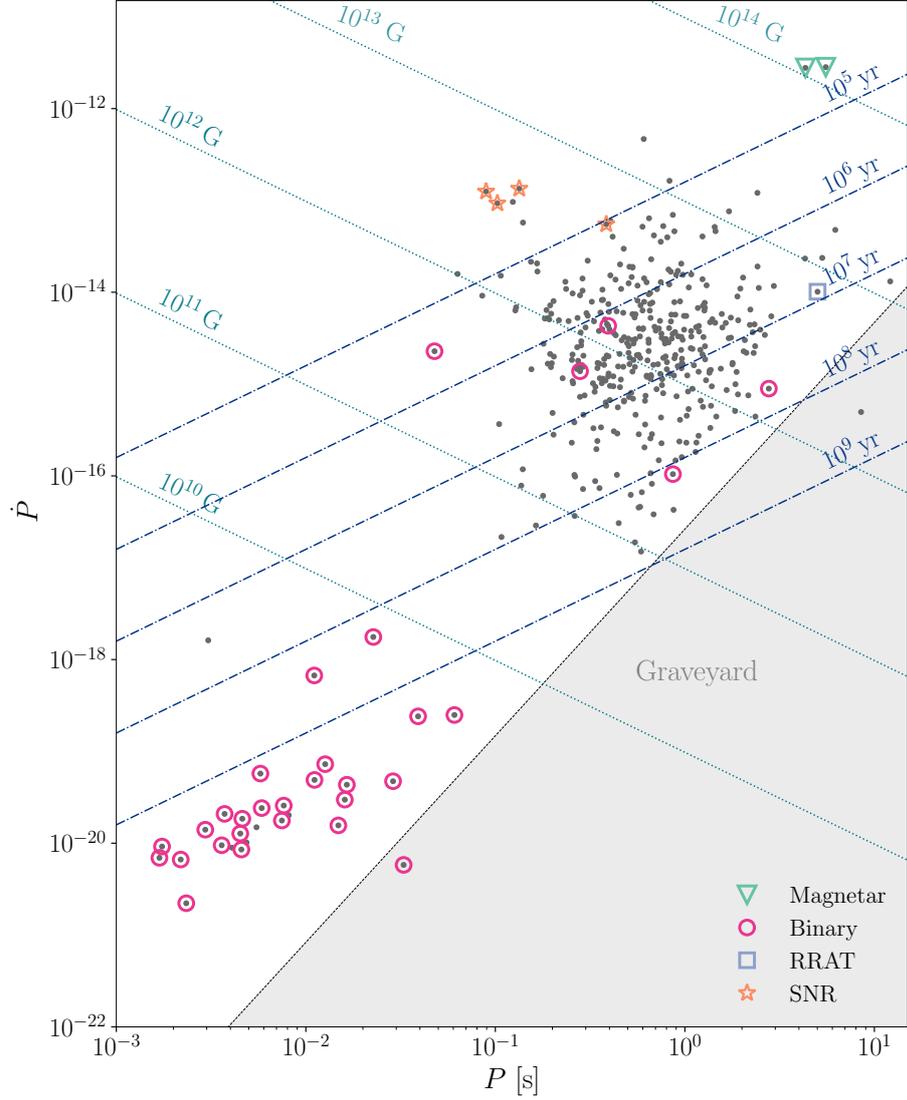


Figure 3.1 Period, period-derivative diagram for all pulsars regularly monitored by UT-MOST. Pulsars residing in binary systems are highlighted by red circles. The RRAT PSR J2033+0042 is indicated by a blue square and the radio-loud magnetars PSR J1622–4950 and XTE J1810–197 by green triangles. Pulsars associated with supernova remnants are highlighted with stars. Lines of constant characteristic age are defined by the dash-dotted lines and constant surface magnetic field strength by dotted lines.

They instead proposed the use of a Bayesian analysis, in which the stochastic timing noise and pulsar properties are modelled simultaneously, avoiding the possibility of biases in the final posterior distributions. This method, in particular the ability to perform model selection, was improved by [Lentati et al. \(2013\)](#) and [Lentati et al. \(2014\)](#), as the [van Haasteren & Levin \(2013\)](#) method is hampered by large matrix inversions and a growing

parameter space as the timing model is linearized.

3.3.1 Phenomenological timing noise model

To model the effects of red timing noise, we utilized the phenomenology outlined in [Lentati et al. \(2014\)](#) and [Lentati et al. \(2016\)](#), where the power spectral density of the red noise process is described in the Fourier domain by a power law⁷ with amplitude A (in units of $\text{yr}^{3/2}$) and spectral index β

$$P_r(f) = \frac{A^2}{12\pi^2} \left(\frac{f}{f_{\text{yr}}} \right)^{-\beta}. \quad (3.1)$$

In addition to the standard power-law model, we also consider a variation of the spectral turnover model of [Coles et al. \(2011\)](#)

$$P_r(f) = \frac{A^2}{12\pi^2} \frac{(f_c/f_{\text{yr}})^{-\beta}}{[1 + (f/f_c)^{-\beta/2}]^2}, \quad (3.2)$$

where f_c is the frequency at which the spectrum turns over. While the models we test are phenomenological rather than drawn from physical theory, models of superfluid turbulence within neutron star interiors ([Melatos & Link, 2014](#)), or the presence of circum-pulsar asteroid belts ([Cordes & Shannon, 2008](#)) predict spectral flattening or turnovers in power spectra of pulsar timing residuals.

Excess white noise in the residuals due to pulse jitter or radio interference can be accounted for by modifying the uncertainties of individual ToAs as

$$\mu_i = \sigma_q^2 + F\sigma_i. \quad (3.3)$$

Here σ_i is the ToA uncertainty on the i -th observation derived from the cross-correlation procedure for generating ToAs, F is a fitting factor (commonly referred to as EFAC) that encodes the contribution of unaccounted instrumental effects and imperfect estimates of ToA uncertainties, while σ_q (error in quadrature: EQUAD) accounts for any additional sources of time-independent uncertainties (e.g pulse jitter).

3.3.2 Bayesian framework

To characterise timing noise and obtain accurate timing models, we used Bayesian parameter estimation to construct posterior probability distributions for the deterministic and

⁷The division by $12\pi^2$ comes from this power-law originally being derived from the one-sided power spectrum a stochastic gravitational-wave background would induce in pulsar timing residuals (e.g [Jenet et al., 2006](#))

Table 3.1 Prior ranges on pulsar and timing noise parameters. Δ_{param} is the uncertainty returned by TEMPO2, T is the length of each pulsar’s data set.

Parameter [units]	Prior range	Prior type
$\nu, \dot{\nu}, \ddot{\nu}$ [Hz, s ⁻² , s ⁻³]	$\pm x^* \times \Delta_{\text{param}}$	Uniform
EFAC	(-1, 2)	Uniform
EQUAD [s]	(-10, 1)	log-Uniform
A [yr ^{3/2}]	(-20, -3)	log-Uniform
β	(0, 20)	Uniform
f_c [yr ⁻¹]	(0.01/ T , 10/ T)	log-Uniform
$\Delta\phi$ [rotations]	(-10, 10)	Uniform
$\Delta\nu_p$ [Hz]	(-12, -5)	log-Uniform
$\Delta\dot{\nu}_p$ [Hz ⁻²]	(-10 ⁻²⁰ , -10 ⁻⁹)	Uniform
$\Delta\nu_d$ [Hz]	(-12, -5)	log-Uniform
τ_d [days]	(0, 3000)	Uniform

* x lies between 100 – 100000 depending on the pulsar.

stochastic pulsar properties θ from the timing residuals r . Prior to performing parameter estimation we first obtained an initial timing solution from previously computed models. In most cases, the initial timing solutions are re-fitted versions of those found in Jankowski et al. (2019) or PSRCAT (Manchester et al., 2005). We fit pulsar parameters that are poorly constrained using a χ^2 minimisation scheme with TEMPO2 (Hobbs et al., 2006; Edwards et al., 2006), resulting in a phase-connected timing solution.

We conducted parameter estimation on these timing models using the TEMPONEST⁸ Bayesian pulsar timing software developed by Lentati et al. (2014). TEMPONEST utilizes the nested sampling algorithm MULTINEST (Skilling, 2004; Feroz & Hobson, 2008; Feroz et al., 2009) to sample the posterior distributions of the parameters θ , given timing residuals r and a timing model \mathcal{M} , while analytically marginalizing over nuisance parameters. The general form of the posterior probability distribution is given by Bayes’ theorem as

$$p(\theta|r, \mathcal{M}) = \frac{\mathcal{L}(r|\theta, \mathcal{M})\pi(\theta, \mathcal{M})}{\mathcal{Z}(r|\mathcal{M})}, \quad (3.4)$$

where $\mathcal{L}(r|\theta, \mathcal{M})$ is the likelihood function for the residuals given a timing model and model parameters (equation 21 of Lentati et al. (2014)), $\pi(\theta, \mathcal{M})$ is our prior knowledge, the ranges of which are listed in Table 3.1 and $\mathcal{Z}(r|\mathcal{M})$ is the Bayesian evidence, which is a single number representing the completely marginalized likelihood defined by

$$\mathcal{Z}(r|\mathcal{M}) = \int d\theta \mathcal{L}(r|\theta, \mathcal{M})\pi(\theta, \mathcal{M}). \quad (3.5)$$

⁸github.com/LindleyLentati/TempoNest

To account for potential covariances when fitting for the parameters of interest, we included the sky-position of the pulsar as free parameters. However, any improvements in the sky-position uncertainty over the values output by TEMPO2 would be marginal at best, as all pulsars in our sample have been timed for more than 1 yr. Hence, we treat the sky-position as a set of nuisance parameters, $\theta_n = \{\alpha, \delta\}$, that are analytically marginalized over to obtain the marginalized posterior distribution for the parameters of interest (θ_i), defined by

$$p(\theta_i|r, \mathcal{M}) = \int \prod_{n \neq i} d\theta_n \pi(\theta_n, \mathcal{M}) \mathcal{L}(r|\theta_i, \theta_n, \mathcal{M}). \quad (3.6)$$

Any glitch parameters in the timing model are also marginalized over, unless we are explicitly attempting to measure them. Neglecting to do so biases the recovered spectral index toward larger values (a steeper red spectrum). For computing posterior distribution confidence intervals we use the maximum likelihood statistics from CHAINCONSUMER⁹ (Hinton, 2016).

After conducting parameter estimation, we can use the resulting Bayesian evidences to compare two or more competing hypotheses ($\mathcal{M}_1, \mathcal{M}_2$) by calculating the odds ratio

$$\mathcal{O}_{12} = \frac{\mathcal{Z}(r|\mathcal{M}_1) \Pi_1}{\mathcal{Z}(r|\mathcal{M}_2) \Pi_2}, \quad (3.7)$$

where Π_1/Π_2 is the *a-priori* odds of the two hypotheses. In our case the prior odds are unity as we assume uninformative priors throughout our analysis. This leaves us with an alternative model comparison metric known as the Bayes factor, which can be calculated as

$$\mathcal{B}_{12} = \frac{\mathcal{Z}(r|\mathcal{M}_1)}{\mathcal{Z}(r|\mathcal{M}_2)} = \frac{\int d\theta_1 \mathcal{L}(r|\theta_1, \mathcal{M}_1) \pi(\theta_1, \mathcal{M}_1)}{\int d\theta_2 \mathcal{L}(r|\theta_2, \mathcal{M}_2) \pi(\theta_2, \mathcal{M}_2)}, \quad (3.8)$$

where θ_1, θ_2 are the parameters associated with models \mathcal{M}_1 and \mathcal{M}_2 respectively. In our analysis, the specific models we compared include:

- White timing noise (WTN): fitting for deterministic pulsar parameters, EFAC and EQUAD only.
- Power-law red noise (PLRN): fitting for a power-law red noise model (equation 3.1) in addition to the WTN parameters.
- Power-law red noise with frequency turnover (PL+FC): includes a turnover in the red power spectrum (equation 3.2) plus WTN parameters.

⁹github.com/samreay/ChainConsumer/

- Second spin-frequency derivative (PLRN+F2): same as PLRN, but also fitting a cubic term to measure $\ddot{\nu}$.

The specific choice of a Bayes factor threshold when performing model selection is largely dependent on what one considers to be an acceptable false positive rate. For instance, a conservative Bayes factor threshold of $|\ln(\mathcal{B}_{12})| > 8$ (corresponding to a false positive rate of $\sim 1/3000$) is generally used in gravitational-wave astronomy (Thrane & Talbot, 2019). A more common interpretation is outlined in Kass & Raftery (1995), where a Bayes factor of $\ln(\mathcal{B}_{12}) > 5$ (false positive rate $\sim 1/150$) is considered to be ‘very strong’ evidence for one hypothesis over the other. In this work we use the latter interpretation, as it has previously been used in pulsar model selection studies (e.g., Lentati & Shannon, 2015; Reardon et al., 2019; Parthasarathy et al., 2019). In cases where neither model is significantly preferred over the other, i.e. for $|\ln(\mathcal{B}_{12})| < 1$, Occam’s razor tells us the least complicated model is preferred.

3.3.3 Braking indices

Over long timescales, the spin-down of a pulsar is often approximated by a power law of the form

$$\dot{\nu} = -K\nu^n, \quad (3.9)$$

where K is a scaling constant related to the pulsar moment of inertia and magnetic field structure (Gunn & Ostriker, 1969) and n is the ‘braking index’. The value of the braking index is potentially an indicator of the physical process that dominates the torque acting to slow the rotation of the neutron star. For instance, a braking index of $n = 1$ arises if the spin-down is dominated by an out-flowing particle wind from the pulsar surface (Harding et al., 1999), $n = 3$ corresponds to magnetic-dipole radiation (e.g. Shapiro & Teukolsky, 1983), and $n = 5$ would indicate the pulsar is spinning down due to some form of quadrupole radiation, such as gravitational waves (Bonazzola & Gourgoulhon, 1996; Yue et al., 2007). Magnetic field evolution or a varying misalignment between the spin and magnetic axes are also predicted to result in $n < 3$ (Blandford & Romani, 1988; Lyne et al., 2013). By taking the derivative of equation 3.9 and solving for n , we can infer the braking index of a pulsar by measuring its *second* spin-frequency derivative ($\ddot{\nu}$), giving

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}. \quad (3.10)$$

Obtaining accurate measurements of $\ddot{\nu}$ is difficult, as measured values of $\ddot{\nu}$ in ‘old’ pulsars are not significantly different from zero. As with measuring ν and $\dot{\nu}$, not accounting for

timing noise in the pulsar residuals when attempting to measure $\dot{\nu}$ will lead to biased measurements, as $\dot{\nu}$ is often highly correlated with timing noise.

3.3.4 Glitch parameter estimation

While the low frequency structures resulting from red noise affect the long term timing precision of pulsars, pulsar glitches result in neutron stars spinning-up on timescales of seconds (Ashton et al., 2019a). This causes a near-instantaneous difference between the observed ToAs and what is expected from the timing model. Some pulsars take days to months to recover toward their original pre-glitch spin frequency (and sometimes do not fully recover or over-recover). In general the change in rotational phase from a glitch can be expressed in terms of instantaneous, permanent changes in the pulsar spin (ν_p) and spin-down ($\dot{\nu}_p$), as well as the exponential spin recovery (ν_d) over time (τ_d)

$$\phi_g(t) = \Delta\phi + \Delta\nu_p(t - t_g) + \frac{1}{2}\Delta\dot{\nu}_p(t - t_g)^2 - \Delta\nu_d\tau_d e^{-(t-t_g)/\tau_d}. \quad (3.11)$$

If the precise epoch at which a glitch occurred (t_g) is poorly constrained then there can be some uncertainty in the precise number of pulsar rotations between the last pre-glitch and first post-glitch observations. Hence an unphysical phase jump ($\Delta\phi$) is frequently implemented to maintain a phase-connected timing solution. If the glitch epoch were known precisely, a phase jump would not be required. Glitch recovery is often associated with the re-pinning of superfluid vortices (Melatos et al., 2008). The degree to which a glitch recovers can be quantified by the recovery parameter $Q = \Delta\nu_d/\Delta\nu_g$, where $\Delta\nu_g = \Delta\nu_p + \Delta\nu_d$.

When fitting for pulsar glitches we include five parameters drawn from equation 3.11 that describe the change in pulsar spin and post-glitch recovery, $\{\Delta\phi, \Delta\nu_p, \Delta\dot{\nu}_p, \Delta\nu_d, \tau_d\}$, in addition to the red noise and spin parameters. For pulsars found to have undergone multiple glitches within our timing data, we fitted all of the relevant glitch parameters simultaneously in order to avoid introducing biases from incomplete glitch models when attempting to model them individually. We then marginalize over the instantaneous phase jump to account for uncertainties on the glitch epoch. We employed Bayesian model selection in cases where it is difficult to tell by eye whether a small glitch or glitch-like event is real, or is the result of a cusp in the residuals due to timing noise.

3.4 Results

3.4.1 Updated timing models

Many pulsars in our sample have improved timing measurements over those from version 1.54 of the ATNF pulsar catalogue¹⁰, including a number that are not present in [Jankowski et al. \(2019\)](#). The updated astrometric and spin parameters are presented in [Appendix A](#). The resulting timing residuals for all pulsars analysed in this work are presented in [Fig. 3.2](#). Ephemeris files in TEMPO2 format, ToAs, plots showing the one- and two-dimensional posterior distributions and a clock correction file are available to download from our online repository¹¹.

3.4.2 Spin frequency second derivatives and braking indices

By simultaneously modelling timing noise as a power-law process, we are able to search for unbiased values of $\ddot{\nu}$ by comparing Bayesian evidences for the power-law red noise (PLRN) and red noise with $\ddot{\nu}$ (PLRN+F2) models. There are 8 pulsars in our sample that marginally prefer the PLRN+F2 model ($1 < \ln(\mathcal{B}) < 5$) for which we recover well constrained $\ddot{\nu}$ posteriors that are inconsistent with zero. We find only one pulsar, PSR J0738–4042, significantly favours the PLRN+F2 model with a log Bayes factor > 5 , while PSR J1001–5507 has a marginally sub-threshold preference ($\ln(\mathcal{B}) = 4.9$). These measurements of $\ddot{\nu}$ along with the inferred braking index for each pulsar are compared with those from PSRCAT in [Table 3.2](#). The sign of the braking indices depends on whether the inferred value of $\ddot{\nu}$ is positive or negative. None of our measurements of $\ddot{\nu}$ and n are consistent with previously published values. Given each of the pulsars with previous $\ddot{\nu}$ measurements have undergone glitches, the difference in results may be due to the accumulated changes in pulsar rotation between measurements or yet to be reported glitches. No glitches have been reported to date in the pulsars that do not have $\ddot{\nu}$ values listed in PSRCAT. While the inferred braking indices are all much larger than the canonical $n = 3$ expected from magnetic dipole radiation, they are consistent with values reported for other young pulsars. It has been speculated these large braking indices may be due to the effects of unmodelled recovery from glitches prior to the start of timing observations ([Johnston & Galloway, 1999](#)). Alternatively, our timing noise model may be incomplete, giving rise to the preference for the PLRN+F2 model. For instance, PSR J0738–4042 may be affected by torque variations initially induced after a profile change in 2005, proposed to be evi-

¹⁰This is the psrscat version from which many of initial ephemerides were drawn from.

¹¹github.com/Molonglo/TimingDataRelease1/

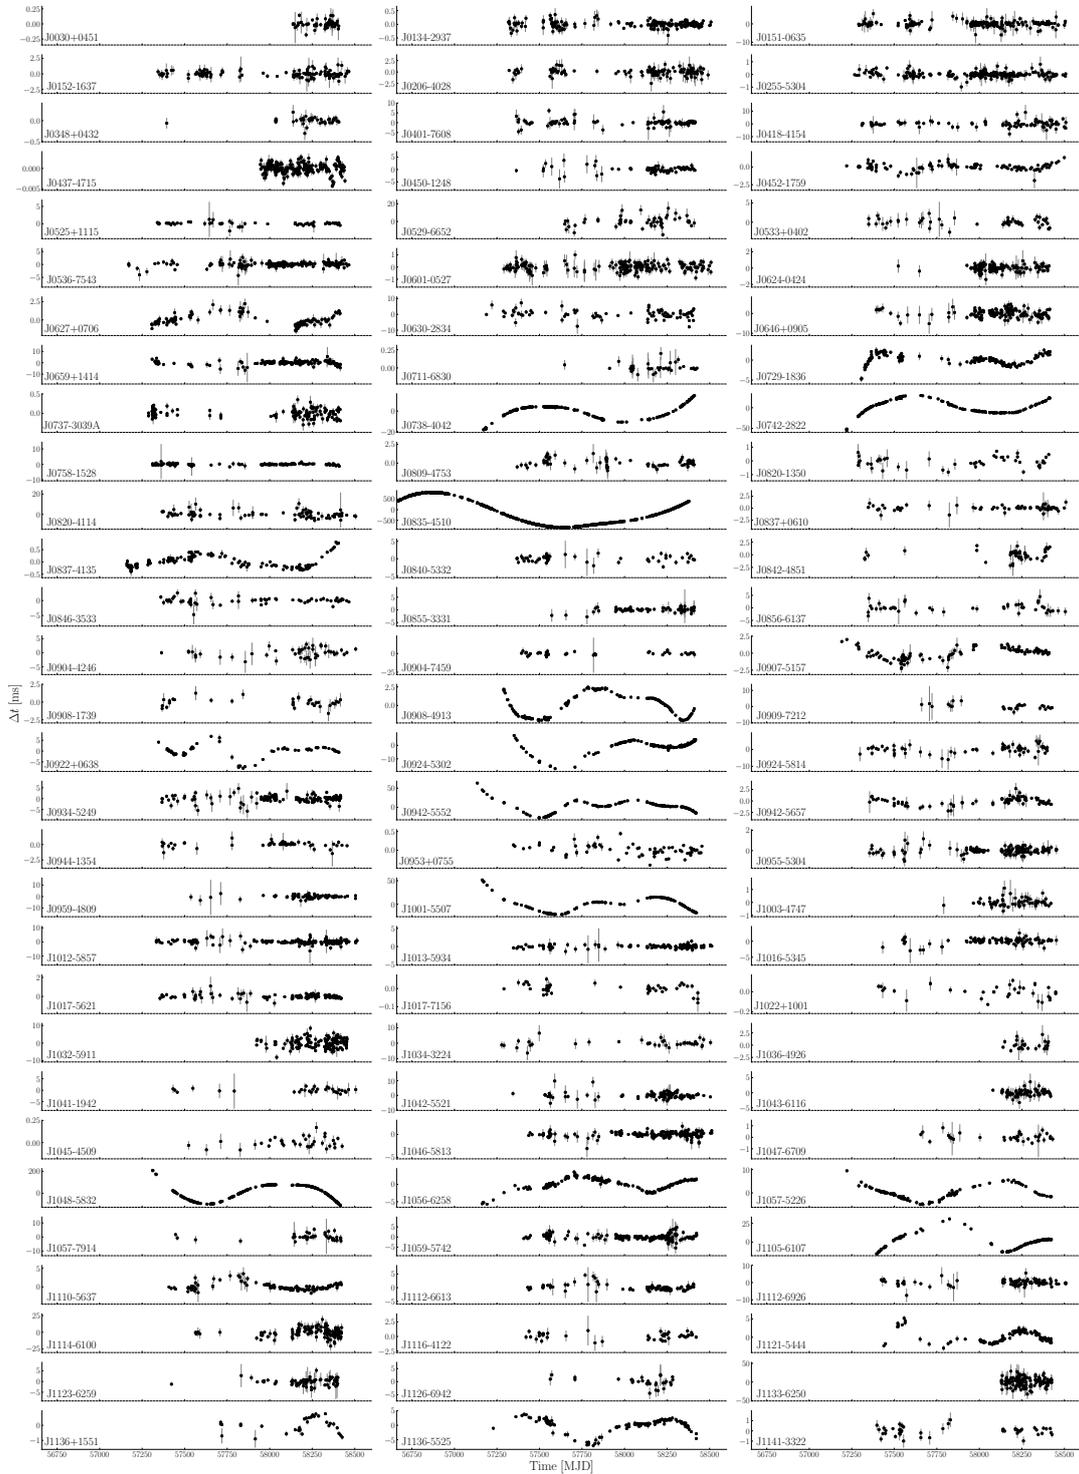


Figure 3.2 Phase connected timing residuals of the pulsars after fitting for ν and $\dot{\nu}$.

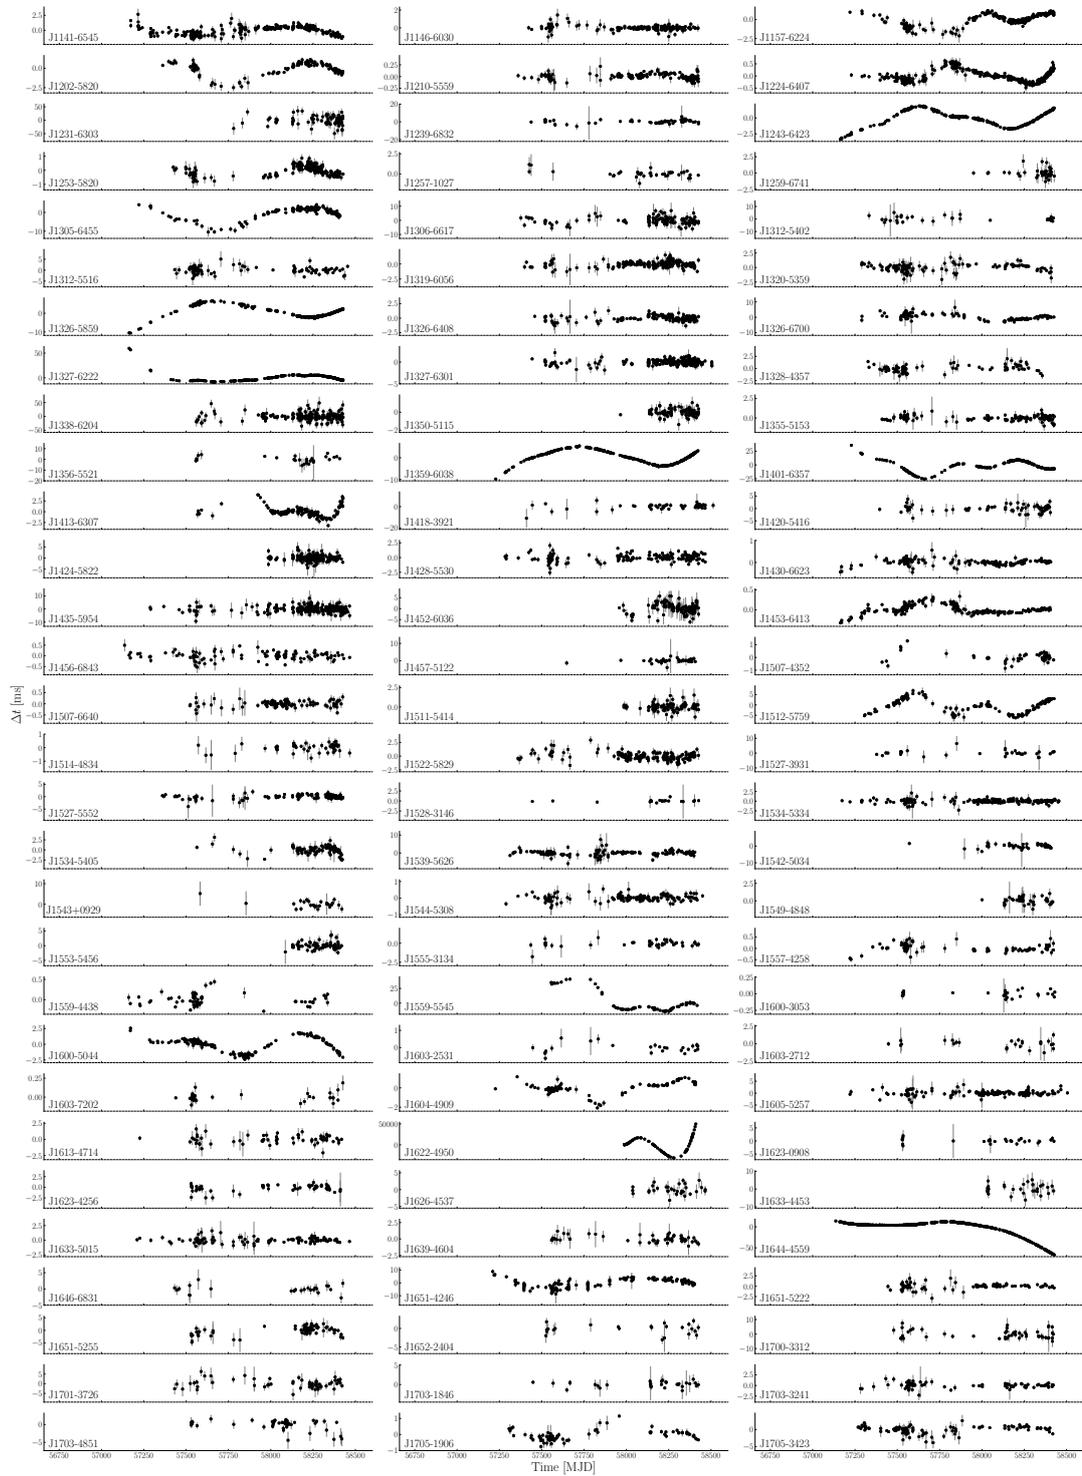


Figure 3.2 (Continued)

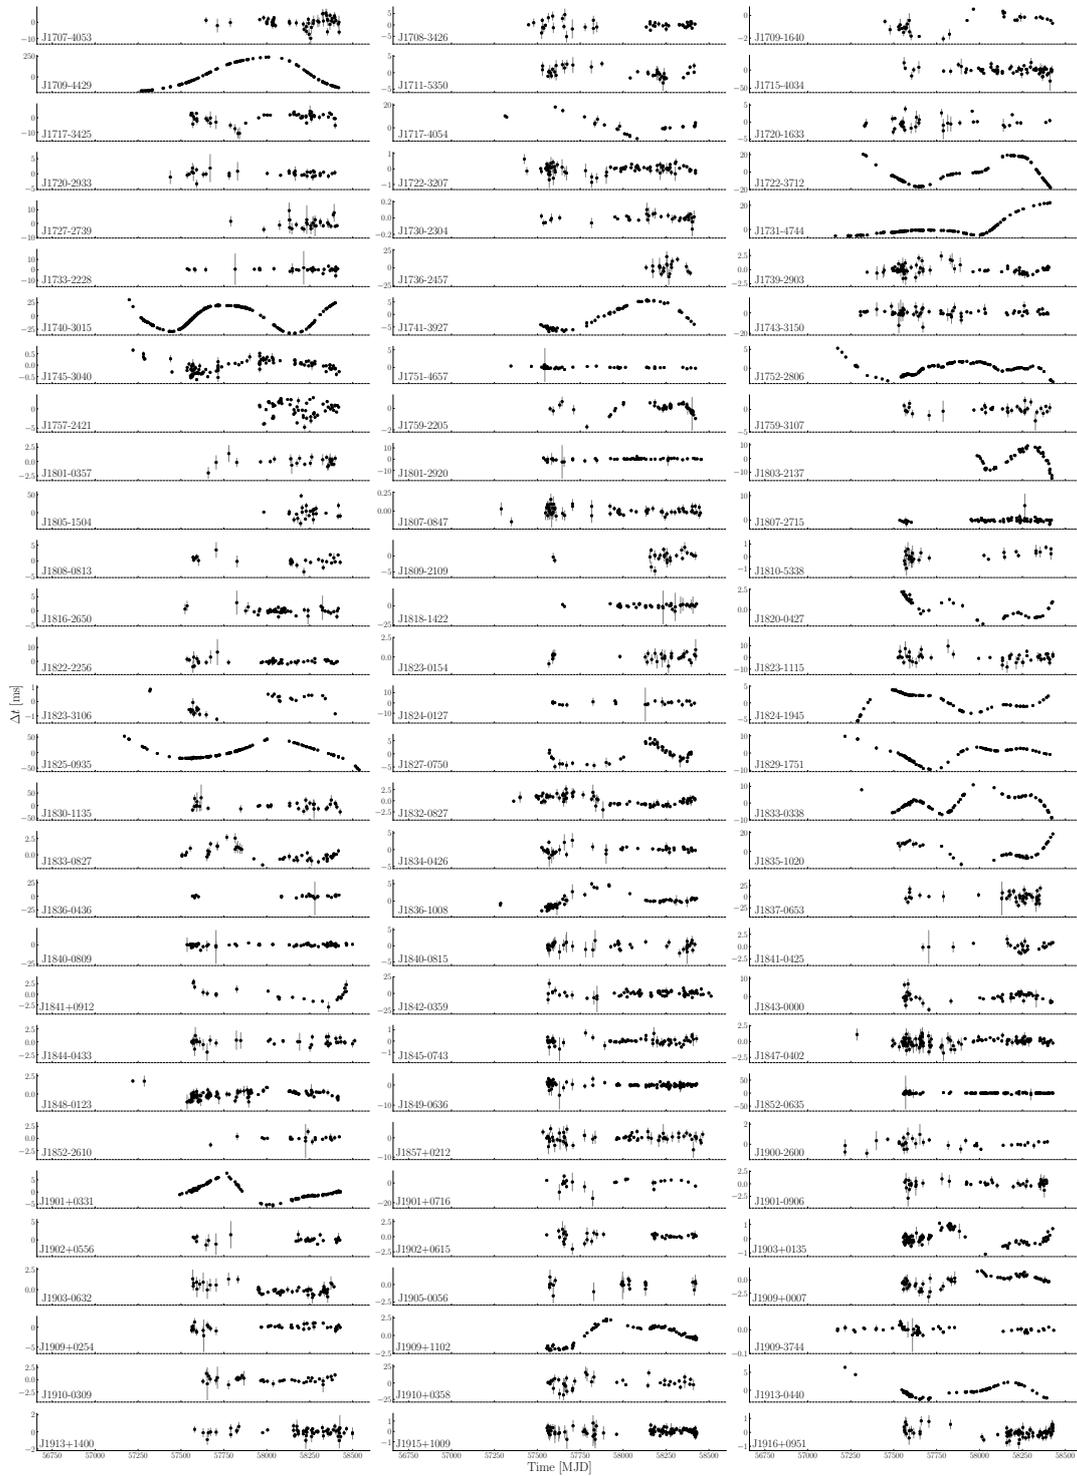


Figure 3.2 (Continued)

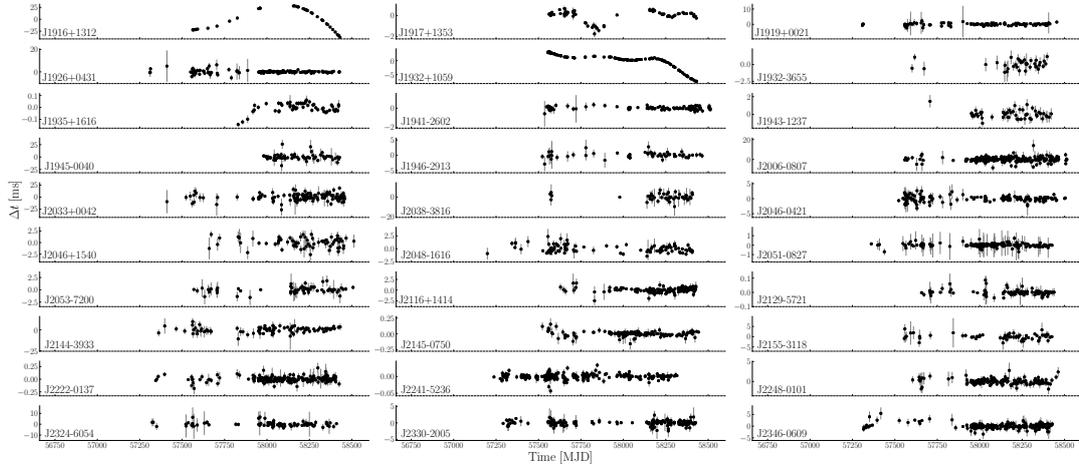


Figure 3.2 (Continued)

dence for an interaction with an asteroid (Brook et al., 2014), while PSR J1001–5507 is known to exhibit discrete spin-down state switching (Chukwude & Buchner, 2012). Neither phenomenon is included in the timing models for these pulsars, and are therefore likely causes of these pulsars’ strong preference for the PLRN+F2 model.

3.4.3 Red timing noise properties

We assessed the presence of red noise in our pulsars using the Bayes factor found from comparing the PLRN model against the WTN model as a detection statistic, where $\ln(\mathcal{B}_{R/W}) > 5$ is a strong detection. Pulsars for which we obtain Bayes factors of $3 < \ln(\mathcal{B}_{R/W}) < 5$ are categorized as ‘probable detections’ since the PLRN model is favoured, but is subject to an increased false positive rate. Those that have Bayes factors in the range $1 < \ln(\mathcal{B}_{R/W}) < 3$ marginally favour a red noise model, but lack the statistical confidence to be distinguishable from the WTN model.

Out of the 300 pulsars analysed, we find 110 strongly favour the PLRN model and 6 that fall into the probable detection category. None of the pulsars in our sample favour the PL+FC model. We find the magnetar PSR J1622–4950 to have the largest red noise amplitude at 1 yr of $\log_{10}(A) = -4.9^{+0.6}_{-0.4}$, with a spectral index $\beta = 7.3^{+3.4}_{-3.6}$. This result should be taken with caution as we do not account for changes in the pulse profile or variations in $\dot{\nu}$ due to short-term changes in magnetic torque that are observed in magnetars (e.g. Camilo et al., 2018). Excluding magnetars, PSR J0835–4510 (the Vela pulsar) has the largest red noise amplitude of any non-recycled pulsar in our data set, with $\log_{10}(A) = -8.2 \pm 0.2$ and a steep spectral index of $\beta = 8.6 \pm 0.9$. This spectral index

Table 3.2 Maximum likelihood posterior measurements of $\ddot{\nu}$ and n compared with reported values in the literature. Values in parentheses indicate the 1σ uncertainties in the last digits, while the errors in our measurements indicate the 95 percent confidence intervals. Only PSR J0738–4042 strongly favours the PLRN+F2 model over the standard PLRN model.

PSR	$\ln(\mathcal{B})$	$\ddot{\nu}$ (10^{-24}s^{-3})	$\ddot{\nu}_{\text{lit}}$ (10^{-24}s^{-3})	n	n_{lit}	τ_c Myr	Glitcher?
J0659+1414	1.5	$1.1^{+1.1}_{-0.5}$	0.764(4)	$21.1^{+2.7}_{-1.2}$	14.44(8)	0.111	Y
J0729–1836	3.6	$-2.4^{+2.3}_{-1.9}$	0.376(15)	$-897.1^{+10.3}_{-5.2}$	139(5)	0.426	Y
J0738–4042	5.4	-3.5 ± 1.2	–	$-96227.0^{+1.8}_{-2.2}$	–	4.32	N
J0942–5552	1.6	$8.1^{+4.2}_{-4.3}$	–	$4591.4^{+3.1}_{-3.5}$	–	0.461	N
J1001–5507	4.9	$1.8^{+0.8}_{-1.5}$	–	$1960.1^{+1.4}_{-4.1}$	–	0.441	N
J1359–6038	4.1	$-3.4^{+3.3}_{-1.1}$	–	$-176.9^{+4.6}_{-1.4}$	–	0.319	N
J1413–6307	1.1	$-4.2^{+3.5}_{-11.2}$	–	$-4606.1^{+2.4}_{-16.5}$	–	0.842	N
J1709–4429	1.0	$106.4^{+45.1}_{-47.3}$	173.1(7)	$13.3^{+0.8}_{-0.3}$	21.35(8)	0.0175	Y
J1909+1102	3.2	$1.1^{+0.7}_{-1.0}$	$-2.02(4)$	$3466.9^{+3.9}_{-5.1}$	$-6615(131)$	1.7	Y

is different to the value measured by [Shannon et al. \(2016\)](#) using 21 years of Vela timing, which may be caused by the occurrence of additional glitches that have recovered since the end of their data set.

We find that three millisecond pulsars favour the PLRN model: PSR J0437–4715, PSR J2145–0750 and PSR J2241–5236. While red noise due to rotational instabilities is known to be present in millisecond pulsars, high precision timing has shown that variations in pulsar DM can mimic timing noise in observations at single frequencies (e.g., [Lentati et al., 2016](#)). Accounting for DM variations requires observing systems that use either wide-band receivers or are capable of observing at multiple frequencies. Due to the limited bandwidth of UTMOST, there is a covariance between DM variations and achromatic timing noise. Hence it is not possible for us to attribute the red noise we observe in millisecond pulsars to rotational irregularities. Unaccounted instrumental artefacts may also contribute to the red noise in these pulsars ([Jankowski et al., 2019](#)).

The lack of multi-band observations also means we cannot infer the contribution of DM variations to the red noise in the non-recycled pulsars. However, [Petroff et al. \(2013\)](#) found only 11 pulsars out of a sample of 160 non-recycled pulsars showed significant changes in DM with time (only setting upper-limits on the remaining 149), while [Shannon et al. \(2016\)](#) showed the Vela pulsar’s DM variations have a sub-dominant contribution to its

overall red noise. Hence, any extra red noise induced by DM variations in our non-recycled pulsar sample would be negligible. Additionally, we find no correlation between red noise parameters and DM.

The full list of the maximum likelihood posterior values and associated 95 percent confidence intervals on the red noise parameters are presented in Appendix A.

3.4.4 Pulsar glitch reanalysis

So far we have observed twelve glitches in eight pulsars, nine of which have been previously reported.¹² The timing residuals for the six pulsars prior to adding glitch corrections are depicted in Fig. 3.3. Cusp-like features in the residuals are the result of large glitches. Note that separate ephemeris and ToA files for PSRs J0835–4510, J1257–1027, J1452–6036 and J1703–4851 that include post-glitch observations and corrections can be found in the online repository¹³. The extended data sets for these pulsars are used only for the glitch analyses, and are not included in our red noise study. Currently pulsars that have undergone a glitch are manually identified in the UTMOST data by searching for glitch-like events in the timing residuals ‘by eye’. This method can be prone to error, with small glitches being glossed over when investigating pulsars that exhibit strong red noise. An automated glitch detection pipeline would be a useful development to search for previously unnoticed glitches in past observations and for near-real time glitch detection.

PSR J0835–4510

There are three reported glitches in PSR J0835–4510 that we have observed with UTMOST (Jankowski et al., 2015b; Palfreyman et al., 2016, 2018; Sarkissian et al., 2019; Kerr, 2019). The first is reported to have occurred on MJD 56922 ± 3 , with a small glitch amplitude of $\Delta\nu_g/\nu = 0.4 \times 10^{-9}$. Our reanalysis returns only upper limits on the change in ν , but does recover $\Delta\dot{\nu}_g = 21.5^{+0.6}_{-1.4} \times 10^{-3}$. However, performing model comparison returns a log Bayes factor of $\ln(\mathcal{B}) = -48.0$, indicating a red noise model without a glitch at this epoch is strongly preferred for this event.

For the second glitch, which was observed in real-time by Palfreyman et al. (2018) at MJD 57734.484991(29), we obtain a glitch amplitude that is similar in magnitude to the published value, while our inferred change in the spin-down of the pulsar is ~ 30 percent smaller than the previously reported value. This is likely due to our analysis incorporating

¹²Seven of these glitches have been added to the Jodrell Bank glitch catalogue (Espinoza et al., 2011a): www.jb.man.ac.uk/pulsar/glitches.html

¹³github.com/Molonglo/TimingDataRelease1/

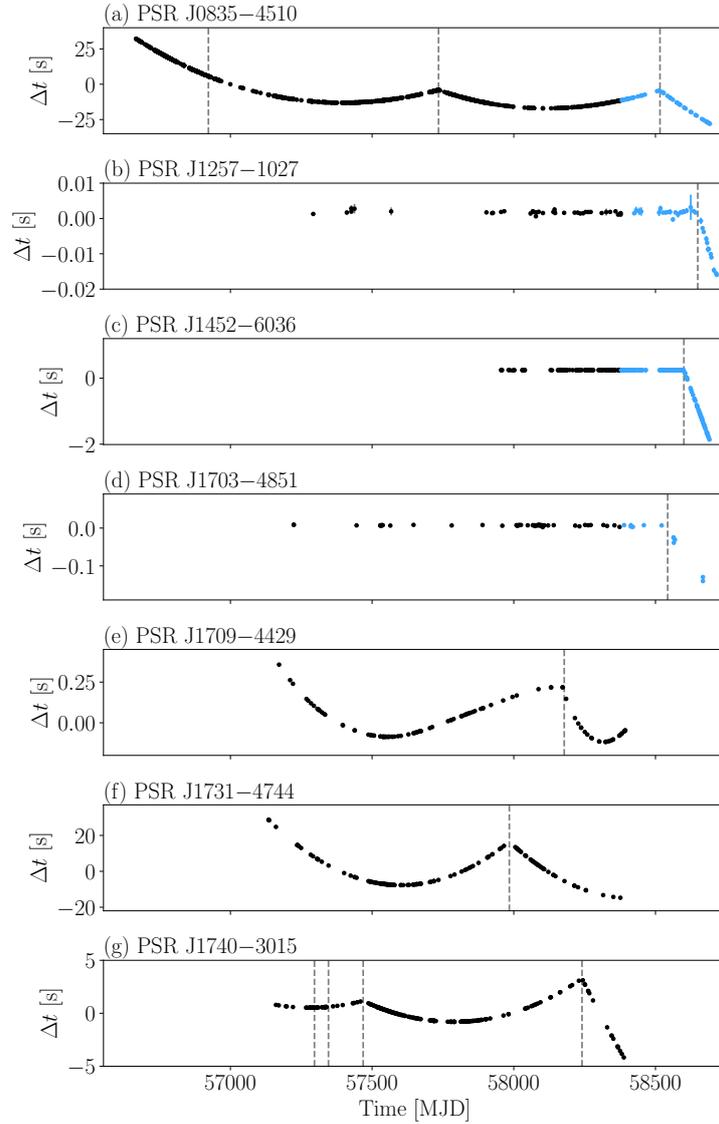


Figure 3.3 Timing residuals for the four glitched pulsars after fitting for ν and $\dot{\nu}$. The epochs of the reported glitches are indicated by the dashed vertical lines. Light blue points in (a), (b), (c) and (d) indicate ToAs that are not used in the red noise analysis.

simultaneous modelling of the glitch and timing noise parameters. Including the short-term recovery found by Sarkissian et al. (2017), we find evidence for additional recovery of $\Delta\nu_d = 4.2^{+0.6}_{-0.3}$ nHz over $12.7^{+3.0}_{-1.2}$ days.

The third glitch occurred during observations by the Hawksbury radio observatory (Sarkissian et al., 2019) and the *Fermi* gamma-ray observatory at MJD 58515.5929(5) (Kerr, 2019). Our recovered glitch amplitude is similar in size to the previously reported values, and is typical of other large Vela glitches ($\Delta\nu_g/\nu \sim 1000$). We also find a small exponential recovery ($Q = 0.005$) occurred over 11.0 ± 1.2 days. We are unable to test whether this glitch underwent any short-term recovery similar to the previous one as our first post-glitch observation was ~ 9 days after the glitch occurred.

In addition to the three glitches we analyse here, a large glitch occurred on MJD 56555.871 (Buchner, 2013; Shannon et al., 2016), prior to the start of our data set. While we cannot constrain the permanent changes in ν and $\dot{\nu}$, we can perform parameter estimation on the recovery parameters in the presence of red noise. We obtain a change in spin frequency of $\Delta\nu_d/\nu = 1591^{+170}_{-322}$, which decays over $\tau_d = 923^{+88}_{-152}$ days. Our measured $\Delta\nu_d/\nu$ is inconsistent with the value inferred by Shannon et al. (2016), but they only have observations up to 116 days after the glitch occurred. Hence any further decay in $\Delta\nu_g$ beyond the end of their data set will have been missed. When compared against the pure PLRN model, the resulting $\ln(\mathcal{B}) = 65.5$ strongly favours the red noise plus glitch recovery model. This confirms that long-term recovery from glitches prior to the start of timing observations can affect the recovered timing parameters.

PSR J1257–1027

This glitch is the first to ever be reported in this pulsar. It is well described by a small permanent change in the pulsar spin ($\Delta\nu_g/\nu = 3.20^{+0.16}_{-0.57} \times 10^{-9}$) with no evidence for recovery. Including a change in the pulsar’s spin-down frequency in our parameter estimation returned only an upper limit of $\Delta\dot{\nu} \lesssim 268 \times 10^{-3}$. Additional observations over longer post-glitch timescales are required for further constraints to be placed on changes in $\dot{\nu}$.

PSR J1452–6036

We discovered a glitch with an amplitude of $\Delta\nu_g/\nu = 270.7^{+0.3}_{-0.4} \times 10^{-9}$ that occurred in PSR J1452–6036 on MJD 58600.29(5). This is the second glitch seen in this pulsar to date and is almost a factor of 10 larger than the glitch observed on MJD 55055.22(4) by Yu et al. (2013). Performing model selection we find a change in spin-down is weakly

Table 3.3 Maximum likelihood posterior values from the glitch parameter posterior distributions and associated 95 percent confidence intervals compared with previously reported measurements. Values in parentheses represent the $1\text{-}\sigma$ uncertainties on the last digit.

PSR	t_g MJD	$\Delta\nu_g/\nu$ $\times 10^{-9}$	$\Delta\dot{\nu}_g/\dot{\nu}$ $\times 10^{-3}$	τ_d days	Q	$(\Delta\nu_g/\nu)_{\text{lit}}$ $\times 10^{-9}$	$(\Delta\dot{\nu}_g/\dot{\nu})_{\text{lit}}$ $\times 10^{-3}$	Ref
J0835–4510	56922(3)	$\lesssim 0.2$	$21.5^{+0.6}_{-1.4}$	–	–	0.4	0.1	1, 2
J0835–4510	57734.484991(29)	$*1448.8^{+0.9}_{-0.8}$	$7.33^{+0.13}_{-0.11}$	$12.7^{+3.0}_{-1.2}$	$*0.011$	1431.24(7)	9.20(83)	3, 4
J0835–4510	58515.5929(5)	$2501.2^{+2.6}_{-3.2}$	$8.69^{+0.28}_{-0.25}$	11.0 ± 1.2	0.005	2491.1(5)	–	5, 6
J1257–1027	58649.3(6)	$3.20^{+0.16}_{-0.57}$	$\lesssim 286$	–	–	–	–	This work
J1452–6036	58600.29(5)	$270.7^{+0.3}_{-0.4}$	$\lesssim 16$	–	–	–	–	This work
J1703–4851	58543.1(3)	$19.0^{+1.0}_{-0.7}$	292^{+38}_{-53}	–	–	–	–	This work
J1709–4429	58178(6)	54.6 ± 1.0	$1.06^{+0.36}_{-0.43}$	$99.1^{+11.3}_{-9.6}$	0.995	52.4(1)	7.30(12)	7
J1731–4744	57984(20)	$3149.5^{+0.5}_{-0.4}$	$1.2^{+0.7}_{-1.1}$	–	–	3147.7(1)	–	8
J1740–3015	57296.5(9)	$0.122^{+0.086}_{-0.081}$	$82.2^{+8.8}_{-8.5}$	–	–	1.30(4)	< 0.66	9, 10
J1740–3015	57346.0(6)	$\lesssim 0.019$	$111.1^{+13.6}_{-8.4}$	–	–	1.94(2)	< 0.07	11
J1740–3015	57468.59(40)	$237.7^{+13.2}_{-9.3}$	$1.71^{+3.24}_{-1.54}$	430^{+91}_{-101}	0.025	229(2)	2.19(4)	11, 10
J1740–3015	58240.781(5)	$842.3^{+7.1}_{-5.6}$	$74.0^{+10.0}_{-13.2}$	–	–	837.88(28)	1.63(14)	12, 10

References indicated in the last column are (1) Jankowski et al. (2015b); (2) Palfreyman et al. (2016); (3) Sarkissian et al. (2017); (4) Palfreyman et al. (2018); (5) Sarkissian et al. (2019); (6) Kerr (2019); (7) Lower et al. (2018); (8) Jankowski et al. (2017); (9) Jankowski et al. (2015a); (10) Espinoza et al. (2011a); (11) Jankowski et al. (2016); (12) Liu et al. (2019). *Includes a short-term $\Delta\nu_d = 129(8)$ nHz recovery over 0.96(17) days (Sarkissian et al., 2017).

disfavoured ($\ln(\mathcal{B}) = 4.2$), hence we can only set an upper limit on $\Delta\dot{\nu}_g/\dot{\nu}$ of $\lesssim 16 \times 10^{-3}$ at the 95 percent confidence level. In addition, we find no evidence for an exponential recovery after this glitch. This could be due to a lack of vortex re-pinning following this glitch, or the recovery having occurred on a timescale too short to be resolved with our current observation cadence (~ 3.3 days between t_g and the first post-glitch observation). Alternatively, the recovery timescale may be significantly longer than our current post-glitch data span.

PSR J1703–4851

The glitch we observed on MJD 58543.1(3) is the first to ever be reported in this pulsar. We recover a moderate change in the pulsar spin of $\Delta\nu_g/\nu = 19.0^{+1.0}_{-0.7} \times 10^{-9}$ and a relatively large change in the spin-down of $\Delta\dot{\nu}_g/\dot{\nu} = 292^{+38}_{-53} \times 10^{-3}$. We find a recovery model is disfavoured for this glitch. While this pulsar is known to undergo emission state switching (Wang et al., 2007) we have only four post-glitch observations to date, hence

we are currently unable to provide any link between the glitch and state switching. The lack of post-glitch observations may also explain why the recovery model is disfavoured, as long-term glitch recoveries require extended observations to detect.

PSR J1709–4429

This glitch is the fourth and smallest glitch observed to date in PSR J1709–4429. The glitch amplitude we recover ($\Delta\nu_g/\nu = 54.6 \pm 1.0 \times 10^{-9}$) is consistent with the previously reported value (Lower et al., 2018), but the change in spin-down frequency was overestimated by a factor of ~ 7 . This is likely due to $\Delta\dot{\nu}$ being covariant with the glitch recovery, which was not fit for in Lower et al. (2018). We find the change-in-spin period due to this glitch almost completely recovers ($Q = 0.995$) in $99.1_{-9.6}^{+11.2}$ days.

PSR J1731–4744

With an amplitude of $\Delta\nu_g/\nu = 3148 \pm 3 \times 10^{-9}$, this is the largest glitch contained within our presented data set, and the largest observed in this pulsar to date (Espinoza et al., 2011a). Previous glitches have shown evidence for linear recoveries (Yu et al., 2013), but we find no evidence for any spin recovery from this glitch.

PSR J1740–3015

Four previously reported glitches have occurred within our timing measurements of PSR J1740–3015 (Jankowski et al., 2015a, 2016; Espinoza et al., 2011a; Liu et al., 2019). We obtain a small amplitude of $\Delta\nu_g/\nu = 0.122_{-0.081}^{+0.086} \times 10^{-9}$ and comparatively large change in spin-down ($\Delta\dot{\nu}_g = 82.2_{-8.5}^{+8.8} \times 10^{-3}$) associated with the first glitch. For the second glitch (MJD 57346.0(0.6)), we are only able to set an upper-limit on the instantaneous change in pulsar spin-frequency. Performing model comparison, we find all models that include the second glitch are strongly disfavoured, suggesting the properties of this glitch are covariant with our red timing noise model. Our analysis of the third glitch recovers a change in spin-frequency that is largely consistent with previously reported values, with a small recovery ($Q = 0.035$) over $430_{-100.9}^{+91.1}$ days. The fourth glitch was discovered in observations of the pulsar at Jodrell Bank (Shaw et al. 2018, private communication, Espinoza et al., 2011a). It was also seen by the Shanghai Tian Ma Radio Telescope (Liu et al., 2019). We find no evidence for spin recovery after this latest glitch. However, the large change in spin-down we recover may be evidence of longer-term recovery, as these two effects are strongly covariant while the pulsar remains in the recovery phase.

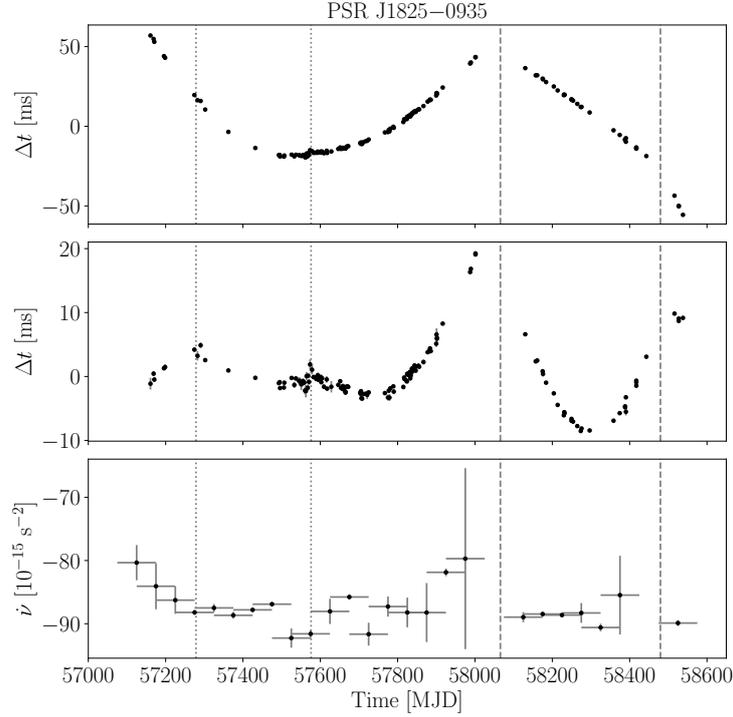


Figure 3.4 Post-fit timing residuals for PSR J1825–0935. Top plot shows the residuals after removing a fit for only ν and $\dot{\nu}$, while the middle plot includes a fit for $\ddot{\nu}$. Bottom plot shows changes in the spin-down frequency over time as determined via fitting $\dot{\nu}$ over ~ 50 day segments (error bars indicate the 1-sigma error). Vertical dashed lines in all three panels correspond to the approximate epochs of the largest events; dotted lines indicate the two smaller events.

Glitch-like events in PSR J1825–0935

Also known as PSR B1822–09, this pulsar has been reported to exhibit timing events, sometimes referred to as ‘slow glitches’ (Zou et al., 2004; Shabanova, 2007). These events are predominantly characterized by a sharp change in the spin-down of the pulsar, which leads to the pulsar spinning up over the course of a few days. This change in spin-down decays exponentially over timescales of days to months. PSR J1825–0935 is also known to switch between two emission states: a ‘B-mode’ where an extra precursor component adjacent to the main pulse is visible, and a ‘Q-mode’, where precursor emission is suppressed and emission from an interpulse component is brightest (Fowler et al., 1981; Morris et al., 1981; Gil et al., 1994). Lyne et al. (2010) showed this switching between these two emission modes is correlated with changes in the spin-down rate, concluding the apparent ‘slow glitches’ are not related to the glitch phenomena, but are instead a result of the pulsar spending more time in one emission/spin-down state versus the other.

Table 3.4 Recovered glitch parameters for the events in PSR J1825–0935.

No.	t_g (MJD)	$\Delta\nu_g/\nu$ ($\times 10^{-9}$)	$\Delta\dot{\nu}_g/\dot{\nu}$ ($\times 10^{-3}$)
1	57278.5(41)	$\lesssim 0.9$	$-1.3^{+2.7}_{-2.0}$
2	57576.1(26)	$\lesssim 0.3$	$-0.9^{+1.1}_{-2.0}$
3	58065.7(3)	$5.2^{+1.7}_{-0.5}$	$-1.1^{+1.7}_{-2.4}$
4	58486.2(9)	$7.6^{+4.5}_{-3.3}$	$14.6^{+37.3}_{-20.0}$

The upper panel of Fig. 3.4 shows the timing residuals for PSR J1825–0935 after subtracting a fit for ν and $\dot{\nu}$. Two glitch-like events are found to have occurred during observing gaps centred at MJD 58065(64) and MJD 58484.8(9)¹⁴. Including a fit for $\ddot{\nu}$ in our timing model reveals two additional events with lower amplitudes that occurred at MJD 57278.5(41) and MJD 57576.1(26). Modelling these four events as standard pulsar glitches, we perform parameter estimation using TEMPO_{NEST} to fit for step changes in ν and $\dot{\nu}$. The recovered glitch parameters are presented in Table 3.4. The first two events are only consistent with upper limits on an instantaneous changes in ν , while the changes in $\dot{\nu}$ both peak at negative values. Changes in $\Delta\dot{\nu}_g$ for all four events are consistent with zero at the 95 percent confidence interval. However, the standard glitch model we employ does not sufficiently describe the true nature of these events. By performing model selection, we find a PLRN model with no glitches is preferred over any glitch+PLRN model, with a $\ln(\mathcal{B}) = 14$ when comparing a PLRN-only model to PLRN+4 glitches and $\ln(\mathcal{B}) = 7.8$ in favour of the PLRN-only model versus a fit to only the two large events. Subtracting off the purely red noise model the post-fit residuals are still dominated by the two larger glitch-like events, implying that at least these two events are not related to the pulsar’s red noise.

The small variations in $\dot{\nu}$ in the bottom panel of Fig. 3.4 at the time of each event are more in line with conventional pulsar glitches as opposed to slow glitches, although the lack of coverage around the more recent large amplitude events means we may have insufficient resolution to detect any rapid changes in spin-down. The mode-changing behaviour and glitch-like events of PSR J1825–0935 demand further investigation, as high-cadence coverage of these events, and any that are discovered in other mode-changing pulsars, may allow us to probe the internal dynamics of these neutron stars.

¹⁴The last event is listed as a glitch in the Jodrell Bank glitch catalogue at MJD 58486.2(9).

3.5 Discussion

3.5.1 Quantifying timing noise strength

While a complete characterisation of pulsar timing noise is yet to be achieved, previous work usually followed one of two approaches. The first involves applying a cubic polynomial to fit for $\ddot{\nu}$ to assess the effects of timing noise. The second uses the root-mean-square (RMS) of the residuals after subtracting a quadratic polynomial, which corresponds to a fit for only ν and $\dot{\nu}$.

Studies undertaken by [Urama et al. \(2006\)](#) and [Chukwude \(2007\)](#) use measurements of $\ddot{\nu}$ to directly infer the strength of the timing noise in their data sets, in addition to searching for correlations with other pulsar parameters. Other users of $\ddot{\nu}$ measurements include work by [Arzoumanian et al. \(1994\)](#) through the use of a model-dependent parameter

$$\Delta_8 = \log \left(\frac{|\ddot{\nu}|}{6\nu} T_8^3 \right), \quad (3.12)$$

where $\ddot{\nu}$ is measured over a total (but arbitrary) observation time span of $T_8 = 10^8$ s. A two-sample variance parameter σ_z is used in [Matsakis et al. \(1997\)](#) to describe pulsar rotational stability

$$\sigma_z = \frac{1}{2\sqrt{5}} \left[\frac{\sigma_{\ddot{\nu}(T)}}{\nu} \right] T^2, \quad (3.13)$$

where $\sigma_{\ddot{\nu}(T)}$ is the RMS of the $\ddot{\nu}$ fit over the observing span T . [Shannon & Cordes \(2010\)](#) note that the Δ_8 method is highly model-dependent since the measured $\ddot{\nu}$ will usually increase on longer timescales, requiring additional time dependent scaling to properly compare values of $\ddot{\nu}$ and Δ_8 . They also state methods based around measurements of $\ddot{\nu}$ (such as σ_z) will often underestimate the amount of timing noise as they neglect contributions from higher-order frequency derivatives.

A method proposed by [Cordes & Helfand \(1980\)](#) assesses the RMS of the total timing noise after conducting a second order fit

$$\sigma_{R,2}^2(T) = \frac{1}{N} \sum_i^N R(t_i)^2, \quad (3.14)$$

where R refers to the timing residuals and N is the number of ToAs. This can be further broken down into red and white components

$$\sigma_{R,2}^2(T) = \sigma_{\text{TN},2}^2(T) + \sigma_{\text{W}}^2(T). \quad (3.15)$$

It is assumed the RMS is usually dominated by $\sigma_{\text{TN},2}^2$ in slow ($P \sim 1$ s) pulsars. The timing noise strength is then estimated via an activity parameter that describes the scaling of $\sigma_{\text{TN},2}$ with respect to PSR J0534+2200 (the Crab pulsar) by

$$A = \log \left[\frac{\sigma_{\text{TN},2}(T)}{\sigma_{\text{TN},2}(T)_{\text{Crab}}} \right]. \quad (3.16)$$

This method assumes the timing noise scales in the same way as the Crab pulsar. [Dewey & Cordes \(1989\)](#), and later [Shannon & Cordes \(2010\)](#) built upon this method by assessing how timing noise varies across the population by examining a scaling relationship between timing noise strength and ν , $\dot{\nu}$ and the observation time span T (in years) as

$$\hat{\sigma}_{\text{TN},2} = C_2 \nu^a |\dot{\nu}|^b T^\gamma. \quad (3.17)$$

Here, the fitting factors C_2 , a , b and γ are measured from the total pulsar population using fits based on maximum likelihood statistics.

We note these methods do not attempt to model the timing noise directly. Instead, they assume the RMS of the residuals accurately describes the timing noise strength. This neglects covariances between intrinsic pulsar properties and red noise, which can result in contaminated residuals as some pulsar properties may be over- or under-fit. This is not an issue for modern Bayesian methods that model both deterministic and stochastic properties simultaneously.

The method we use for assessing timing noise strength was developed in parallel with [Parthasarathy et al. \(2019\)](#), in which timing noise strength is inferred from the red noise amplitude and spectral index, obtained via parameter estimation with `TEMPONEST`, and the observation span as

$$\sigma_{\text{RN}}^2 = A^2 T^{\beta-1}. \quad (3.18)$$

Using this metric, we find the magnetar PSR J1622–4950 has the strongest timing noise in our sample. However, as stated earlier the torque variations due to the magnetar’s decaying magnetic field, rather than spin noise, are expected to dominate the observed red noise. In addition, [Shannon & Cordes \(2010\)](#) argued that timing noise in magnetars is statistically different to that in millisecond and non-recycled pulsars. Given the red noise we observe in both the millisecond pulsars and PSR J1622–4950 can be explained via processes other than rotational irregularities, we restrict our analysis to the 280 non-recycled pulsars in our sample, 112 of which strongly favour the PLRN model.

3.5.2 Correlations with individual pulsar properties

A number of pulsar properties can be inferred from their spin and spin-down. These include the characteristic age (τ_c), surface dipole magnetic field strength (B_{surf}) and the rotational kinetic energy loss over time (\dot{E}). Simplified approximations to these properties, along with the pulse period derivative (\dot{P}), can be expressed in terms of the spin and spin-down frequencies, as follows

$$\begin{aligned}\dot{P} &\propto \nu^{-2}|\dot{\nu}| \\ \tau_c &\propto \nu|\dot{\nu}|^{-1} \\ B_{\text{surf}} &\propto \nu^{-3/2}|\dot{\nu}|^{1/2} \\ \dot{E} &\propto \nu|\dot{\nu}|.\end{aligned}\tag{3.19}$$

Similar to previous work on pulsar timing noise, we examine whether correlations exist between the measured red noise strength and these pulsar properties, in addition to the spin and spin-down on their own. Assuming timing noise strength scales with spin and spin-down frequencies in a similar fashion to Equation 3.17, we compared the inferred strength against a predictive metric

$$\chi_{\text{RN}} \propto \nu^a |\dot{\nu}|^b,\tag{3.20}$$

where the values of a and b can be set to the approximate pulsar properties we are comparing.

From a frequentist perspective, the amount of correlation between σ_{RN} and χ_{RN} can be quantified via the Pearson correlation coefficient

$$r_p = \frac{\sum_{i=1}^N (\sigma_{\text{RN},i} - \mu_\sigma)(\chi_{\text{RN},i} - \mu_\chi)}{[\sum_{i=1}^N (\sigma_{\text{RN},i} - \mu_\sigma)^2 \sum_{i=1}^N (\chi_{\text{RN},i} - \mu_\chi)^2]^{1/2}},\tag{3.21}$$

where $\mu_\sigma = \frac{1}{N} \sum_{i=1}^N \sigma_{\text{RN},i}$ is the mean of the σ_{RN} values, and μ_χ is the mean of χ_{RN} . However, this approach does not take into account potential covariances between the means of $\sigma_{\text{RN},i}$ and $\nu^a |\dot{\nu}|^b$, or scatter in the measurements. It is also not robust against the influence of outliers in the data set.

An alternative approach involves assuming the red noise measurements and values generated from equation 3.20 are correlated samples drawn from an underlying bivariate Gaussian distribution, the shape of which is best described by the set of hyper-parameters

Table 3.5 Frequentist (r_p) and Bayesian (ρ) correlations between pulsar properties and red noise strength, using only strong red noise detections (D) and including lower confidence detections (D + PD). Errors represent the 95 percent confidence intervals.

	D		D + PD	
	r_p	ρ	r_p	ρ
ν	0.13	$0.18^{+0.29}_{-0.26}$	0.16	$-0.18^{+0.33}_{-0.25}$
$\dot{\nu}$	0.45	$0.47^{+0.29}_{-0.20}$	0.46	$0.47^{+0.27}_{-0.20}$
\dot{P}	0.49	$0.51^{+0.28}_{-0.18}$	0.49	$0.48^{+0.24}_{-0.21}$
τ	-0.51	$-0.51^{+0.19}_{-0.25}$	-0.51	$-0.53^{+0.17}_{-0.26}$
B_{surf}	0.37	$0.39^{+0.30}_{-0.22}$	0.37	$0.38^{+0.28}_{-0.22}$
\dot{E}	0.39	$0.41^{+0.30}_{-0.21}$	0.41	$0.42^{+0.29}_{-0.20}$

$\{\mu_\sigma, \mu_\chi, \sigma_\sigma, \sigma_\chi, \rho\}$ as

$$\mathcal{N}_2(\sigma_{\text{RN}}, \chi_{\text{RN}}) = \frac{1}{2\pi\sigma_\sigma\sigma_\chi\sqrt{1-\rho^2}} \exp\left[\frac{-1}{2(1-\rho^2)} \times \left(\frac{\sigma_{\text{RN}}^2}{\sigma_\sigma} + \frac{\chi_{\text{RN}}^2}{\sigma_\chi} - \frac{2\rho\sigma_{\text{RN}}\chi_{\text{RN}}}{\sigma_\sigma\sigma_\chi}\right)\right]. \quad (3.22)$$

Here, $\mu_\sigma, \sigma_\sigma$ are the mean and variance of the distribution in the σ_{RN} direction, and μ_χ, σ_χ represent the mean and width in the χ_{RN} direction. The parameter ρ indicates the direction in which the bivariate Gaussian is rotated and provides an estimate for the level of correlation between σ_{RN} and χ_{RN} . For simplicity, we express the bivariate Gaussian as $\mathcal{N}_2(\sigma_{\text{RN}}, \chi_{\text{RN}}) = \theta^T \mathbf{C}^{-1} \theta$, where \mathbf{C} is the covariance matrix

$$\mathbf{C} = \begin{bmatrix} \sigma_\sigma^2 & \rho\sigma_\sigma\sigma_\chi \\ \rho\sigma_\sigma\sigma_\chi & \sigma_\chi^2 \end{bmatrix}, \quad (3.23)$$

and $\theta = (\sigma_{\text{RN}} - \mu_\sigma, \chi_{\text{RN}} - \mu_\chi)$. We can then write the likelihood function from which our samples are drawn from as

$$\mathcal{L}(\theta|\mathbf{C}) = \frac{1}{2\pi\sqrt{|\mathbf{C}|}} \prod_{i=1}^N \exp\left[\frac{-1}{2}\theta_i^T \mathbf{C}^{-1} \theta_i\right], \quad (3.24)$$

where $|\mathbf{C}| = \sigma_\sigma^2\sigma_\chi^2(1-\rho^2)$. We use the BILBY software library (Ashton et al., 2019b) and PYMULTINEST (Buchner et al., 2014), a Python wrapper for the MULTINEST algorithm, to sample the hyper-parameter posterior distributions using the 112 non-recycled pulsars that strongly favour the PLRN model, ignoring those that are consistent with the WTN model. The resulting Frequentist and Bayesian correlation coefficients for pulsars with

strong evidence for red noise (D; $\ln(\mathcal{B}) > 5$) and when including those with less confident evidence (D + PD; $3 < \ln(\mathcal{B}) < 5$) are presented in Table 3.5.

We find the strongest correlations exist with \dot{P} and $\dot{\nu}$, in addition to a similar anti-correlation with characteristic age. These correlations are smaller than those from the σ_z analysis performed by Hobbs et al. (2010), who presented an analysis of the ongoing timing campaign of a large sample of pulsars at the Jodrell Bank Observatory ($N = 366$, $T_{\text{mean}} \sim 19$ yr), but are similar to those from Namkham et al. (2019) who assessed the timing noise of 129 ‘middle-aged’ ($\tau_c \sim 1$ Myr) pulsars observed by the Parkes radio telescope over ~ 4 yr using the σ_z metric. We find pulsar spin-frequency has effectively no correlation with timing noise, but the weak correlation of 0.3 from Hobbs et al. (2010) does overlap with the 95 percent confidence region of our Bayesian correlation parameter for the pulsars that strongly prefer the PLRN model. These differences are to be expected as Hobbs et al. (2010) included both millisecond and partially-recycled pulsars when calculating their correlation coefficients, while we are limited to non-recycled pulsars.

3.5.3 Scaling relation fitting and hyper-parameter estimation

To see how the timing noise strength varies independently of specific pulsar properties, we build upon previous work by Dewey & Cordes (1989) and (Shannon & Cordes, 2010) for finding a scaling relation that maps timing noise strength to a scaling of pulsar ν and $\dot{\nu}$, and observing timescale T . We rewrite their scaling relation (Equation 3.20) as

$$\chi_{\text{RN}} = \xi \nu^a |\dot{\nu}|^b T^\gamma, \quad (3.25)$$

where Shannon & Cordes (2010) use the symbol C_2 in place of ξ to represent the linear scaling factor. Unlike in Section 3.5.2, a and b are not set to fixed values to approximate certain pulsar properties. Instead we allow them to be free parameters with uniformly sampled priors ($-10 < \pi(a, b) < 10$).

Shannon & Cordes (2010) define the joint likelihood function

$$\mathcal{L}(\sigma_{\text{RN}}, \sigma_{\text{UL}} | \chi_{\text{RN}}, \varepsilon) = \mathcal{L}(\sigma_{\text{RN}} | \chi_{\text{RN}}, \varepsilon) \times \mathcal{L}(\sigma_{\text{UL}} | \chi_{\text{RN}}, \varepsilon), \quad (3.26)$$

which is comprised of a standard Gaussian likelihood with an additional hyper-parameter ε^2 to describe the scatter in the measured red noise strength in the pulsar sample, which may be attributed to variations in the amount of turbulence in their superfluid interiors (Melatos & Link, 2014), and an upper-limit likelihood. The likelihood functions

contributing to Equation 3.26 are given by

$$\mathcal{L}(\sigma_{\text{RN}}|\chi_{\text{RN}}, \varepsilon) = \prod_i^N \frac{1}{\sqrt{2\pi\varepsilon^2}} \exp\left[-\frac{(\sigma_{\text{RN},i} - \chi_{\text{RN},i})^2}{2\varepsilon^2}\right], \quad (3.27)$$

and

$$\mathcal{L}(\sigma_{\text{UL}}|\chi_{\text{RN}}, \varepsilon) = \prod_j^N 1 - \frac{1}{2} \operatorname{erfc}\left[-\frac{(\sigma_{\text{UL},j} - \mu_{\text{RN},j})}{\varepsilon\sqrt{2}}\right], \quad (3.28)$$

where σ_{UL} is the ‘upper limit’ on the red noise strength and erfc is the complementary error function. We use this likelihood to calculate posterior distributions for the scaling hyper-parameters as follows

$$p(\chi_{\text{RN}}, \varepsilon|\sigma_{\text{RN}}, \sigma_{\text{UL}}) \propto \mathcal{L}(\sigma_{\text{RN}}, \sigma_{\text{UL}}|\chi_{\text{RN}}, \varepsilon)\pi(\chi_{\text{RN}}, \varepsilon). \quad (3.29)$$

This modified version of the [Shannon & Cordes \(2010\)](#) formalism (mSC10 hereafter) requires making two key assumptions: scatter in the maximum likelihood posterior values of σ_{RN} due to measurement uncertainties are either negligible or absorbed by ε , and the upper limit likelihood holds true for the pulsars with only a marginal preference (i.e. $1 < \ln(\mathcal{B}) < 3$) for the PLRN model. It also does not take into account the information that can be gained by including the full posterior distribution for σ_{RN} during the fitting.

We can overcome these shortcomings by assuming our measurements of σ_{RN} for a given pulsar is drawn from a Gaussian distribution, the mean of which depends on the aforementioned scaling of the pulsars spin and spin-down frequencies (equation 3.25), and a variance ε^2 defined by

$$\pi(\sigma_{\text{RN}}|\chi_{\text{RN}}, \varepsilon) = \frac{1}{\sqrt{2\pi\varepsilon^2}} \exp\left[-\frac{(\sigma_{\text{RN}} - \chi_{\text{RN}})^2}{2\varepsilon^2}\right]. \quad (3.30)$$

This distribution represents an approximation to the ‘true’ probability distribution of σ_{RN} across the population. To compute the posterior distributions for our scaling hyper-parameters ($\{\xi, a, b, \gamma\}$), we use the marginalized likelihood

$$\mathcal{L}(\mathbf{r}|\chi_{\text{RN}}, \varepsilon) = \int d\sigma_{\text{RN}} \mathcal{L}(\mathbf{r}|\sigma_{\text{RN}})\pi(\sigma_{\text{RN}}|\chi_{\text{RN}}, \varepsilon). \quad (3.31)$$

As we are using an ensemble of N individual pulsars with residuals $\mathbf{r} = \{r_1, \dots, r_N\}$, we can take the product of the individual likelihoods to obtain the total likelihood of the timing

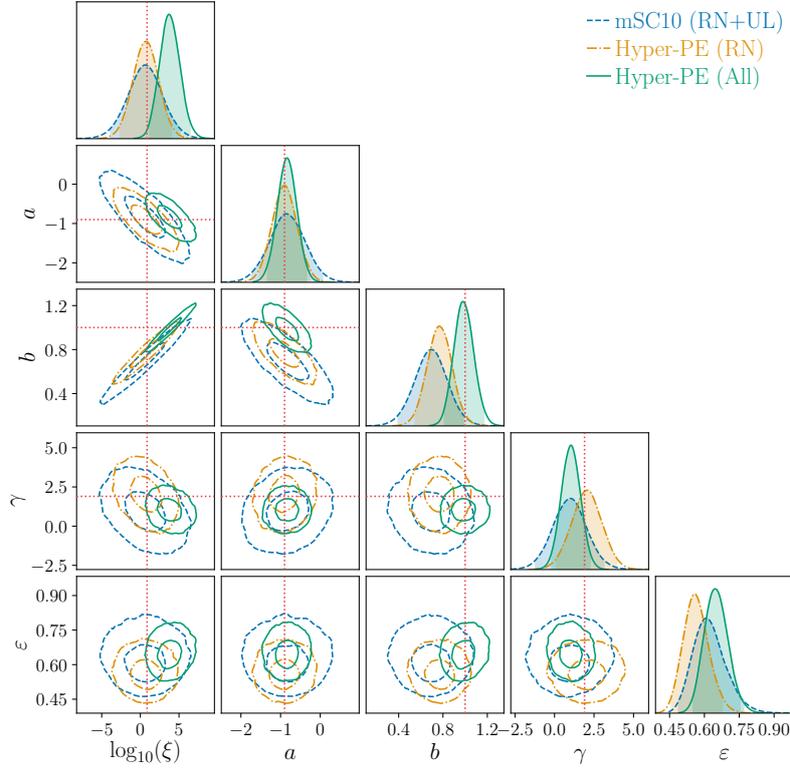


Figure 3.5 One- and two-dimensional posterior distributions for the scaling parameters across the non-recycled pulsar population. Contours in the two-dimensional posteriors indicate the 50 and 95 percent confidence regions. Shading in the one-dimensional posteriors covers the 95 percent confidence intervals. The red dotted lines indicate the non-recycled pulsar maximum likelihood values from [Shannon & Cordes \(2010\)](#).

data given the red noise strength for the population

$$\mathcal{L}_{\text{tot}}(\mathbf{r}|\sigma_{\text{RN}}) = \prod_i^N \mathcal{L}(r_i|\sigma_{\text{RN},i}), \quad (3.32)$$

hence the total marginalized likelihood can be rewritten as

$$\mathcal{L}_{\text{tot}}(\mathbf{r}|\chi_{\text{RN}}, \varepsilon) = \prod_i^N \int d\sigma_{\text{RN},i} \mathcal{L}(r_i|\sigma_{\text{RN},i}) \pi(\sigma_{\text{RN},i}|\chi_{\text{RN},i}, \varepsilon). \quad (3.33)$$

From Bayes theorem, we can find $\mathcal{L}(r_i|\sigma_{\text{RN},i})$ as

$$\mathcal{L}(r_i|\sigma_{\text{RN},i}) = \mathcal{Z}(r_i) \frac{p(\sigma_{\text{RN},i}|r_i)}{\pi(\sigma_{\text{RN},i})}, \quad (3.34)$$

where the prior on $\sigma_{\text{RN},i}$ is the product of the log-uniform prior on A and the uniform

Table 3.6 Comparison between recovered maximum likelihood posterior values for the scaling hyper-parameters and their associated 95 percent confidence intervals.

Parameter	mSC10 (RN+UL)	Hyper-PE (RN)	Hyper-PE (All)
$\log_{10}(\xi)$	$0.6^{+4.3}_{-4.4}$	$1.0^{+3.4}_{-3.0}$	$3.7^{+2.4}_{-2.7}$
a	$-0.87^{+0.83}_{-0.91}$	$-0.88^{+0.63}_{-0.60}$	$-0.84^{+0.47}_{-0.49}$
b	0.69 ± 0.29	$0.77^{+0.21}_{-0.22}$	$0.97^{+0.16}_{-0.19}$
γ	1.00 ± 2.17	$2.1^{+2.0}_{-1.8}$	1.0 ± 1.2
ε	$0.60^{+0.11}_{-0.16}$	$0.56^{+0.10}_{-0.12}$	$0.64^{+0.11}_{-0.16}$

prior on β

$$\pi(\sigma_{\text{RN}}) = \pi(A, \beta) = \pi(A)\pi(\beta) = \frac{1}{A}, \quad (3.35)$$

for $A \in \{10^{-20}, 10^{-3}\}$. We then substitute equation 3.34 into equation 3.33 to obtain

$$\mathcal{L}(\mathbf{r}|\chi_{\text{RN}}, \varepsilon) = \int d\sigma_{\text{RN},i} \mathcal{Z}(r_i) p(\sigma_{\text{RN},i}|r_i) \frac{\pi(\sigma_{\text{RN},i}|\chi_{\text{RN},i}, \varepsilon)}{\pi(\sigma_{\text{RN},i})}. \quad (3.36)$$

Converting the integral over $\sigma_{\text{RN},i}$ to a sum over the posterior samples, as $\int dx p(x)f(x) \approx 1/n_s \sum_i^{n_s} f(x_i)$ (Hogg & Foreman-Mackey, 2018), the final likelihood function is

$$\mathcal{L}(\mathbf{r}|\chi_{\text{RN}}, \varepsilon) = \prod_i^N \frac{\mathcal{Z}(r_i)}{n_i} \sum_k^{n_i} \frac{\pi(\sigma_{\text{RN},i,k}|\chi_{\text{RN},i}, \varepsilon)}{\pi(\sigma_{\text{RN},i,k})}, \quad (3.37)$$

where k is the number of posterior samples for the i -th pulsar. We can then combine this likelihood with the prior for the hyper-parameters, $\pi(\chi_{\text{RN}}, \varepsilon)$, and the Bayesian evidence for the timing data, to obtain the posterior distributions for the hyper-parameters

$$p(\chi_{\text{RN}}, \varepsilon|\mathbf{r}) = \frac{\mathcal{L}(\mathbf{r}|\chi_{\text{RN}}, \varepsilon)\pi(\chi_{\text{RN}}, \varepsilon)}{\mathcal{Z}(\mathbf{r})}. \quad (3.38)$$

As with equation 3.29, the posterior distributions are sampled using PYMULTINEST. We present the resulting one- and two-dimensional posterior distributions in Fig. 3.5, comparing results from the mSC10 method, hyper-parameter estimation (Hyper-PE) using only the 112 non-recycled pulsars that favour the PLRN model, and the resulting improvement when all 280 non-recycled pulsars are included in the Hyper-PE method regardless of the preferred model.

Recovered values for each scaling hyper-parameter from both methods are listed in Table 3.6. It is clear the Hyper-PE method returns improved estimates over the mSC10

method (with the exception of ε , which is consistent between all three methods), as indicated by the smaller confidence regions. Including the additional 168 white noise dominated pulsars provides additional improvements, as the Hyper-PE method takes into account additional information by summing over the entire posterior distribution of σ_{RN} , rather than only using the maximum likelihood posterior value. Our recovered value of ε differs from the value of $\varepsilon = 1.6 \pm 0.1$ reported by [Shannon & Cordes \(2010\)](#). This inconsistency could be due to the use of two different methods of modelling timing noise in pulsars, resulting in a different amount of measurement scatter.

We can compare our results to those in the literature by looking at the specific scaling relation from equation 3.20. From our Hyper-PE method, we find the timing noise strength of the non-recycled pulsars in our sample follow the scaling relation

$$\chi_{\text{RN}} \propto \nu^{-0.84^{+0.47}_{-0.49}} |\dot{\nu}|^{0.97^{+0.16}_{-0.19}}. \quad (3.39)$$

[Shannon & Cordes \(2010\)](#) computed a scaling relation of $\sigma_{\text{TN},2} \propto \nu^{-0.9 \pm 0.2} |\dot{\nu}|^{1.0 \pm 0.05}$, or $\sigma_{R,2} \propto \nu^{-0.7 \pm 0.1} |\dot{\nu}|^{0.76 \pm 0.02}$ when including the effects of additional white noise, while the analysis by [Hobbs et al. \(2010\)](#) found the relation $\sigma_z(10 \text{ yr}) = 10^{-11.5} \nu^{-0.4} |\dot{\nu}_{-15}|^{0.8}$, where $|\dot{\nu}_{-15}|$ is the spin-down rate in units of 10^{-15} s^{-2} . More recently, [Parthasarathy et al. \(2019\)](#) made use of TEMPONEST, and the same timing noise strength metric we used, to analyse 85 ‘young’ ($\tau_c \lesssim 1 \text{ Myr}$), high- \dot{E} pulsars with ~ 10 years of timing observations. Using a grid search to find the maximally correlated ν scaling index – at a fixed scaling parameter of 1 for $\dot{\nu}$ – they found a scaling relation of $\sigma_P \propto \nu^{-0.9 \pm 0.1} |\dot{\nu}|^1$. This same grid search method was also used by [Namkham et al. \(2019\)](#) to infer their scaling of the σ_z parameter, obtaining the relation $\sigma_P \propto \nu^{-1.7} |\dot{\nu}|^{1.0}$.

The values of a and b from each of these relations are compared with our results in Fig. 3.6. Our relation is entirely consistent with [Shannon & Cordes \(2010\)](#)’s $\sigma_{\text{TN},2}$ scaling, while both relations from [Hobbs et al. \(2010\)](#) and [Parthasarathy et al. \(2019\)](#) fall within our 95-percent confidence regions. Improving our measurements of a and b can be achieved by adding additional pulsars to our sample and/or by extending the lengths of our timing baselines. The improvement made by adding more pulsars is illustrated by the ~ 22 percent reduction in the Hyper-PE confidence regions in Fig. 3.5 after including the 168 pulsars that favour the WTN model in our analysis. Additional observations over longer timing baselines may allow us to obtain improved red noise amplitude and spectral index measurements, and detect low amplitude red noise in pulsars that currently favour the WTN model.

Applying a consistent approach to measuring timing noise strength in various data

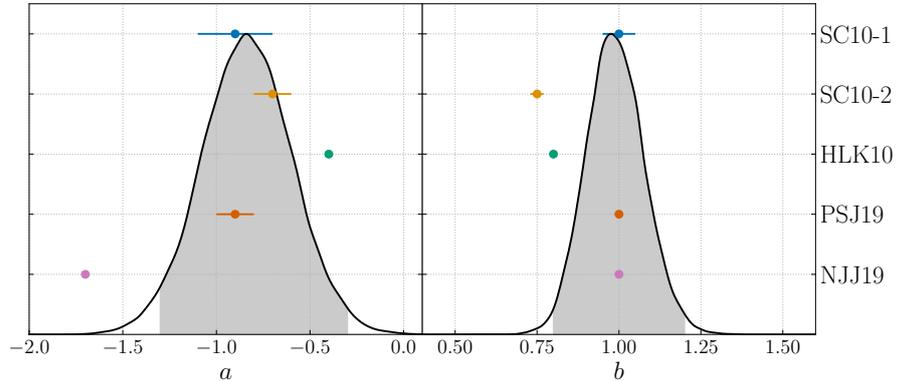


Figure 3.6 Comparison between our Hyper-PE (All) posteriors for a and b (black curves, shading represents the 95 percent confidence regions), and values (with $1\text{-}\sigma$ errors) from [Shannon & Cordes \(2010\)](#) (SC10-1: $\sigma_{\text{TN},2}$ and SC10-2: $\sigma_{R,2}$), [Hobbs et al. \(2010\)](#) (HLK10), [Parthasarathy et al. \(2019\)](#) (PSJ19) and [Namkham et al. \(2019\)](#) (NJJ19).

sets is of particular importance when it comes to comparing observations with theoretical models of timing noise processes. Our method of performing parameter estimation on the stochastic properties of individual pulsars with `TEMPONEST` followed by using hyperparameter estimation to infer the scaling across the population can be easily extended to other large pulsar timing programmes, or even modified to accommodate astrophysically motivated distributions on the expected spectral properties of timing noise (see, e.g. [Melatos & Link, 2014](#)). Model selection studies could also allow for different physical timing noise models to be compared, along with their implications for our understanding of the dynamic processes and internal structure of neutron stars.

3.5.4 Two noteworthy pulsars

Here we discuss results for two pulsars of particular interest: PSR J0737–3039A, for which we constrain the decay of its orbital period due to gravitational-wave emission, and PSR J1402–5124, whose celestial coordinates we find to be different to published values.

PSR J0737–3039A

J0737–3039A is the ‘A’ pulsar of the renowned double pulsar system discovered by [Burgay et al. \(2003\)](#). The ‘B’ pulsar ([Lyne et al., 2004](#)) is currently not visible due to its magnetic-axis precessing out of our line-of-sight ([Perera et al., 2010](#)). As its name suggests, the double pulsar allowed a determination of the mass ratio R by measuring the two semi-major axes of the pulsars. When combined with the sum of the masses derived from

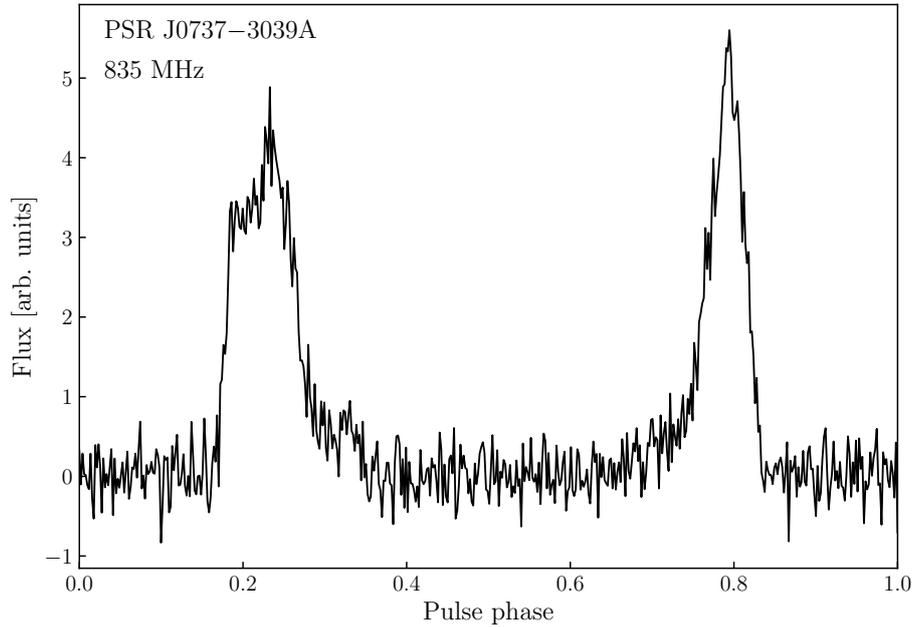


Figure 3.7 Average pulse profile of PSR J0737–3039A at 835 MHz after summing 71.1 hours worth of observations taken over four years.

the advance of periastron, this completely determines the constituent masses to high precision, and predicts the rate of orbital decay due to the emission of gravitational waves. Using `TEMPONEST` to conduct parameter estimation on the pulsar’s rotational and binary parameters, we find the relativistic properties of the system to be consistent with the masses and GR parameters measured by [Kramer et al. \(2006b\)](#). This produces the integrated profile seen in Fig. 3.7. The (albeit limited) timing precision is good enough for us to spot any potential glitches in the pulsar’s rotation, assist in dispersion measure variation monitoring, and to be used in undergraduate projects to demonstrate post-Keplerian effects such as advance of periastron and orbital decay to better than 1 percent accuracy.

PSR J1402–5124

During our regular FRB-search transit observations, the real-time detection pipeline reported a candidate pulse from an unknown source with a DM of 53 pc cm^{-3} and a S/N of 10.2 at MJD 58657.40992245. Upon inspection of the data, we detected many faint single pulses with similar morphology drifting through our fan-beams. A periodicity search on the data ‘stitched’ according to the sky-drift-rate revealed a high S/N pulsar candidate with a period of 1.38 s and $\text{DM} = 51(9) \text{ pc cm}^{-3}$, closely matching the properties of pulsar PSR J1402–5124 reported by [Manchester et al. \(1978\)](#). A first-order localisation of

the source, however, yielded a sky position that was inconsistent with the coordinates reported in the pulsar catalogue. Tracking the source using finely-spaced fan-beams over the next few days, we optimized the coordinates of the pulsar to: RA = 14:02:56.0(2), DEC = -50:21:43(49). The improved declination measurement is consistent with the value of DEC = -50:20(5) reported by [Edwards et al. \(2001\)](#). In [Fig. 3.8](#) we highlight the variability of the pulse profile by plotting the phase vs time of the pulsar throughout a 40 minute observation after placing a tied-array-beam on the updated coordinates. Using 5 epochs of timing observations, we are able to constrain the spin-period of the pulsar to $P = 1.380182295(4)$ s. Subtracting this new period measurement from the value reported in [Manchester et al. \(1978\)](#) we derive an estimated spin-down of $\dot{P} = -5.413(4) \times 10^{-15}$, placing it in the population of ‘middle-aged’ pulsars ($\tau_c \approx 4$ Myr). The astrometric and rotational properties will be further constrained as we continue to time the pulsar.

3.6 Conclusions

We have performed an initial study of the rotational properties of 300 bright, southern-sky radio pulsars observed by UTMOST using the Bayesian pulsar timing software `TEMPONEST` to characterise the stochastic properties of our pulsar sample and to obtain unbiased measurements of ν and $\dot{\nu}$. Three millisecond pulsars in our sample favour the Power-Law Red Noise (PLRN) model, but this excess noise is due to a combination of instrumental artefacts and extrinsic astrophysical sources as opposed to rotational irregularities.

We also used `TEMPONEST` to reanalyse nine previously reported pulsar glitches. While the posterior distributions we recover for the change in spin-frequency are generally consistent with previously published values, we are able to place tighter constraints on the change in spin-down and spin-recovery. Performing model selection, we find a PLRN-only model is preferred for two previously reported glitches, one in PSR J0835–4510 on MJD 56922(3) and PSR J1740–3015 on MJD 57346.0(6). This highlights the importance of accounting for timing noise of a pulsar when estimating glitch properties, and the potential use of model selection as a means of verifying glitch detections any additional, undiscovered glitch/micro-glitch candidates that may be present in the data. It also indicates conducting reliable parameter estimation on micro-glitch candidates in the presence of strong red noise is difficult. We also present the discovery of a new glitch in PSR J1452–6036 and the first ever glitch observed in PSR J1703–4851. Additionally, we observed four unusual glitch-like events in PSR J1825–0935, the larger of which cannot be adequately explained by timing noise alone. While it is possible the two large events were due to ‘slow glitches’, we have insufficient observations around their epochs to obtain high resolution measure-

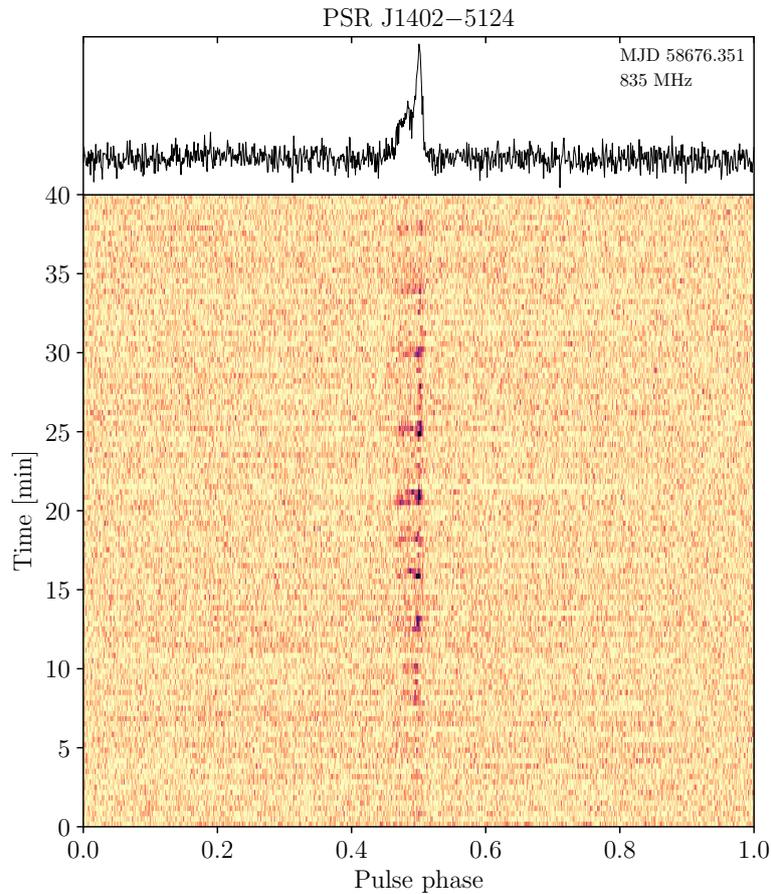


Figure 3.8 Stack of 20 s sub-integrations over 40 minutes (bottom) and the integrated pulse profile (top) of PSR J1402–5124 at 835 MHz. Brightening (dimming) of the pulsar toward the lower (upper) edge of the figure is due to the telescope’s beam response. Dynamic range has been reduced to highlight profile changes between sub-integrations.

ments of $\dot{\nu}$. As a result, we do not observe a sharp increase in the spin-down typically associated with slow glitches.

Limiting ourselves to only the non-recycled pulsars in our sample, we find the strongest correlation between various pulsar properties and the relative red noise strength exists with pulsar spin-down (\dot{P}) and $\dot{\nu}$, with a similar anti-correlation with characteristic age. These correlations agree with recent work by [Namkham et al. \(2019\)](#), but are generally weaker than those found by [Hobbs et al. \(2010\)](#). We conclude this difference is likely caused by [Hobbs et al. \(2010\)](#) including millisecond pulsars in their analysis, while we are limited to non-recycled pulsars. Building on work by [Shannon & Cordes \(2010\)](#), we developed a new Bayesian hyper-parameter estimation (Hyper-PE) framework for inferring the scaling between red noise strength and pulsar spin-frequency and spin-down across

the population. This relation can be used to estimate the expected red noise strength of a pulsar based on its measured spin and spin-down. Our inferred scaling relation of $\chi_{\text{RN}} \propto \nu^a |\dot{\nu}|^b$, where $a = -0.84_{-0.49}^{+0.47}$ and $b = 0.97_{-0.19}^{+0.16}$, is consistent with those found in previous studies by [Shannon & Cordes \(2010\)](#), [Hobbs et al. \(2010\)](#) and [Parthasarathy et al. \(2019\)](#). As timing noise strength scales with the length of each pulsar data span, applying our Hyper-PE method to a much larger data set with longer timing baselines will enable more stringent constraints to be placed on the scaling between timing noise strength and pulsar rotational properties. These data could be obtained by UTMOST in the future, or other large, long-term timing programmes such as those undertaken at Jodrell Bank and CHIME/Pulsar ([Ng, 2018](#)). A natural extension of our study would be to include measurements of red noise in a large sample of millisecond pulsars and magnetars. In addition, the ability to perform model selection studies using astrophysically motivated noise models could allow us to place constraints on the precise mechanism behind spin noise.

Finally, we discussed the capability of UTMOST to contribute to the monitoring of relativistic binary systems such as the double pulsar PSR J0737–3039. We also used the interferometer nature of the instrument to measure an updated position for the bright, mode-changing pulsar PSR J1402–5124 in addition to providing the first estimate of this pulsar’s spin-down rate.

Addendum

It was later shown by [Dunn et al. \(2021\)](#) that our recovery of the glitch parameters in PSR J1709–4429 was affected by a previously unknown degeneracy between $\Delta\nu_g$ our observing cadence. The observations of this pulsar by UTMOST are spaced by almost an integer number of days, which introduces an ambiguity in the precise pulse-numbering assigned to the ToAs. This results in multiple possible glitch solutions existing. Our subsequent re-analysis of this glitch in Chapter 4 was unaffected by this issue owing to the (largely) non-periodic scheduling of Parkes observations.

4

The impact of glitches on the rotational evolution of young pulsars

We report on a timing programme of 74 young pulsars that have been observed by the Parkes 64-m radio telescope over the past decade. Using modern Bayesian timing techniques, we have measured the properties of 124 glitches in 52 of these pulsars, of which 74 are new. We demonstrate that the glitch sample is complete to fractional increases in spin-frequency greater than $\Delta\nu_g^{90\%}/\nu \approx 8.1 \times 10^{-9}$. We measure values of the braking index, n , in 33 pulsars. In most of these pulsars, their rotational evolution is dominated by episodes of spin-down with $n > 10$, punctuated by step changes in the spin-down rate at the time of a large glitch. The step changes are such that, averaged over the glitches, the long-term n is small. We find a near one-to-one relationship between the inter-glitch value of n and the change in spin-down of the previous glitch divided by the inter-glitch time interval. We discuss the results in the context of a range of physical models.

This chapter has been published in full in *Monthly Notices of the Royal Astronomical Society*.

4.1 Introduction

Pulsars are highly magnetised, rotating neutron stars that are observed as sources of electromagnetic pulses with a periodicity that matches their rotation rates. The spin frequency, ν , of a pulsar decreases over time as it loses rotational kinetic energy through various processes. This phenomenon can be described in terms of ν and the spin-down rate ($\dot{\nu}$) as a simple power-law of the form

$$\dot{\nu} = -\kappa\nu^n. \tag{4.1}$$

Here the coefficient κ depends on the magnetic dipole moment amplitude, angle between magnetic and spin axes and moment of inertia of the neutron star, while n is the braking index of the dominant torque acting on the pulsar over time. Under the assumption that κ remains constant over time, the braking index can be found by re-arranging the time derivative of Equation 4.1 such that

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}, \quad (4.2)$$

where $\ddot{\nu}$ is the second derivative of pulsar spin frequency. Braking indices are often measured by either performing local fits to ν and $\dot{\nu}$ across many years of observations and fitting the slope of the resulting $\dot{\nu}$ measurements (e.g. Lyne et al., 1996), or through long-term coherent pulsar timing with a single model (e.g. Parthasarathy et al., 2019). In the latter method, the rotation phase of a pulsar over time is described as a truncated Taylor series

$$\phi(t) = \phi_0 + \nu(t - t_0) + \frac{1}{2}\dot{\nu}(t - t_0)^2 + \frac{1}{6}\ddot{\nu}(t - t_0)^3 + \dots, \quad (4.3)$$

where ϕ_0 is the rotation phase at some reference time t_0 , and the braking index is inferred from measurements of ν and its derivatives via Equation 4.2. Approximating a pulsar as a rotating dipole in a vacuum results in the star spinning down purely through dipole radiation with a braking index of $n = 3$ (e.g. Shapiro & Teukolsky, 1983). However, pulsars are neither simple bar magnets, nor do they exist in a vacuum. Their magnetospheres are filled with energetic plasma, some of which is accelerated away from the neutron star on open field lines as powerful particle winds. Spin-down dominated by such outflows would result in a braking index satisfying $n \approx 1$ for the hypothetical case of a split-monopole outflow (Michel & Tucker, 1969), and $n \lesssim 3$ for a co-rotating magnetosphere modelled as an extended dipole (Melatos, 1997) or a superposed vacuum and force-free structure (Contopoulos & Spitkovsky, 2006). It has also been suggested that pulsars spinning down due to gravitational-wave radiation may have $n = 5$ (Bonazzola & Gourgoulhon, 1996; Yue et al., 2007) or $n = 7$ (Owen et al., 1998; Alford & Schwenzer, 2014). None of these models take into account the interior structure of pulsars or the coupling between the crust of the star and its magnetosphere, which almost certainly affect the observed rotational-evolution of neutron stars.

A small sample of young (characteristic age, $\tau_c < 30$ kyr) pulsars have been found to possess values of n between these theoretical values and are claimed to represent the long-term rotational evolution of these pulsars on timescales that are much longer than their inter-glitch intervals (see Tables 1 and 4 of Espinoza et al. 2017). However, many pulsars

have measured $\ddot{\nu}$, and hence n , that span a large range of values (e.g. [Namkham et al., 2019](#)). These are often referred to as being ‘anomalous’ in the literature as they deviate significantly from the small values expected from radiative mechanisms. Additionally, many of these measurements can be attributed to various stochastic processes arising from either the magnetosphere or internal dynamics. Hence, the large inferred values of n do not necessarily reflect the true long-term rotational evolution of these pulsars over many decades. For example, the braking indices inferred from simple fits of Equation 4.3 to the arrival times of a large samples of pulsars without accounting for these stochastic processes tend to be almost equally split between positive and negative values, induced by the presence of low-frequency stochastic variations in the pulse arrival times, often referred to as timing noise ([Hobbs et al., 2010](#)). Glitches, sudden spin-up events that can be caused by some form of stress build-up and release process (see [Haskell & Melatos 2015](#) for a review of glitch mechanisms), can result in discontinuities in the otherwise smooth spin-down of pulsars. Many glitching pulsars have been found to exhibit steep, positive gradients in $\dot{\nu}$ (i.e. a large $\ddot{\nu}$) in-between subsequent glitches (e.g. [Yu et al., 2013](#)). Unlike the braking indices inferred from arbitrary cubic fits, these large ‘inter-glitch’ braking indices are consistently found with values of n between ~ 10 -200 ([Johnston & Galloway, 1999](#)). They are often associated with a particular form of post-glitch behaviour, namely large step-changes in $\dot{\nu}$ followed by an ‘linear’ $\dot{\nu}$ recovery, which was first noticed in the timing of PSR J0835–4510 (B0833–45: the Vela pulsar, [Cordes et al. 1988](#)). This phenomenon has been interpreted as possible evidence for the creeping of thermally unpinned superfluid vortices between pinning sites inside neutron stars ([Alpar et al., 1984b,a, 1993](#); [Haskell et al., 2020](#)), though such theoretical models are often difficult to falsify (see Section 7 of [Haskell & Melatos 2015](#)).

However, recent work by [Parthasarathy et al. \(2019\)](#) challenges the assertions that large n must result from either stochastic or glitch-based processes. Using a modern Bayesian inference framework, they discovered a sample of 19 young pulsars with high spin-down energies (\dot{E}), that possess large, predominately positive n despite accounting for various types of timing noise. The robustness of these measurements were reinforced in a follow-on study ([Parthasarathy et al., 2020](#)), where timing models that included long-term exponential glitch recoveries were found to be inconsistent with the data, while the inclusion of up to a decade of additional timing data had little effect on the recovered values of n . They further demonstrated these measurements could not be explained by the presence of unaccounted glitch recovery effects. The lack of observed glitches in these pulsars indicates that the underlying mechanism responsible for the large n must be stable

over decade-long timescales.

One idea to explain these large values of n is to relax the assumption that κ in Equation 4.1 remains constant with time. Physically, this could correspond to one of (or potentially a combination of) e.g. changes in the neutron star moment of inertia over time (Ho & Andersson, 2012), evolution of the magnetic and spin axes towards (or away from) alignment (Goldreich, 1970; Tauris & Manchester, 1998; Melatos, 2000), or changes in the surface magnetic field strength (Viganò et al., 2013; Ho et al., 2015). Observationally, variations in κ with time manifest themselves in the braking index itself evolving on \sim kyr timescales, a process that has been exploited in some population synthesis studies to explain the broad distribution pulsar of spin periods and period-derivatives (e.g. Johnston & Karastergiou, 2017). Hence the robust measurement of n of a large sample of pulsars could allow us to place constraints on the mechanisms responsible for the long-term rotational evolution of pulsars.

For this work, we analysed a group of 74 young pulsars that have been observed over the last decade as part of the young pulsar timing programme on the Parkes radio telescope. We derive the parameters for 124 glitches in 52 of these pulsars, and combine our measurements with those of Parthasarathy et al. (2019, 2020) to explore the rotational evolution of the pulsars. The structure of the paper is as follows. In Section 4.2, we briefly outline the observing and data processing steps, while the glitch search and inference frameworks that we employed are detailed in Section 4.3. We present a timing noise limited catalogue of 124 pulsar glitches from our pulsar sample as well as our inferred upper-limits on the minimum glitch size across the sample and statistical analyses of the overall glitch properties in Section 4.4. We also briefly outline the updated pulsar properties obtained as part of our model selection studies. In Section 4.5 we highlight our braking index measurements for 33 pulsars and explore differences in the implied versus observed long-term rotational evolution of these pulsars. Conclusions and potential future directions are summarized in Section 4.6.

4.2 Observations

The young pulsar timing project (P574) has been running on the CSIRO Parkes 64-m radio telescope (also known as *Murriryang*) with an approximately monthly cadence since the beginning of 2007. Originally intended to aid in the detection of pulsed gamma-ray emission with the *Fermi* satellite’s Large Area Telescope (Smith et al., 2008; Weltevrede et al., 2010), the sample has changed somewhat over the years and now consists of some 260 pulsars (Johnston et al., 2021). For the purposes of this paper we do not re-analyse

the timing of those pulsars already reported in Parthasarathy et al. (2019, 2020), nor those pulsars that were added to the project in 2014 (Namkham et al., 2019). This leaves a total of 74 objects, mainly pulsars with a high spin-down energy, $\dot{E} \gtrsim 10^{34}$ ergs s⁻¹.

The pulsars in our sample were observed in the 20-cm band using the multi-beam, H-OH and UWL receivers (Staveley-Smith et al., 1996; Granet et al., 2011; Hobbs et al., 2020). All observations were folded in real-time using the polyphase digital filterbank signal processors to form PSRFITS format archive files (Hotan et al., 2004), each with 1024 phase bins and 1024 frequency channels covering 256 MHz of bandwidth. Each folded archive was then excised of radio frequency interference, before being flux and polarization calibrated using the tools in PSRCHIVE (Hotan et al., 2004; van Straten et al., 2012). After averaging the individual observations in time, frequency and polarization to obtain a one-dimensional profile of flux versus pulse phase, pulse times of arrival (ToAs) are then generated by cross-correlating the averaged pulse profiles in the Fourier-domain with a smoothed, high S/N template (Taylor, 1992). A more comprehensive description of the observations and data processing can be found in Johnston & Kerr (2018) and Johnston et al. (2021).

We also make use of extended data sets available for 27 pulsars from the Parkes Observatory Pulsar Data Archive (Hobbs et al., 2011) that we pre-pended to the beginning of the P574 data. These ‘legacy’ data comprise observations undertaken prior to 2007 using the multi-beam and H-OH receivers at 20-cm wavelengths with the analogue and digital signal processors as presented in Wang et al. (2007) and Yu et al. (2013).

4.3 Pulsar inference framework

In order to measure the properties of pulsars we must first obtain (approximately) phase connected timing solutions. For young pulsars this is not always a simple task as the presence of timing noise and glitches within the timing data often result in a loss of phase coherence in the pulse arrival times. Hence a complete catalogue of glitches in our pulsars is needed in order to characterise their timing properties.

4.3.1 Glitch detection and phase connection

For a large fraction of the pulsars in our sample, we were able to identify when a glitch had occurred via visual inspection of their timing residuals. Preliminary fitting to the glitches using TEMPO2 (Hobbs et al., 2006; Edwards et al., 2006) allowed us to assign pulse numbers to each ToA – i.e. determine the integer number of rotations that have

occurred since the first observation. However, in cases where a particularly large glitch had occurred, obtaining a preliminary fit became impossible. To solve this problem we employed two separate methods: estimates of change in spin and spin-down from local measurements of ν and $\dot{\nu}$ obtained from stride-fits to a moving window containing 5-6 ToAs, and the Hidden Markov Model (HMM) glitch detection algorithm developed by Melatos et al. (2020). Local ν - $\dot{\nu}$ measurements allowed us to obtain a rough estimate of the glitch size that was needed to obtain both a coherent solution and verify the correct pulse numbering was applied. The HMM algorithm provided both an independent means of obtaining preliminary measurements of the permanent step changes in spin frequency ($\Delta\nu_p$) and spin-down frequency ($\Delta\dot{\nu}_p$) associated with a glitch, and an automated means for detecting any additional glitches that were missed by visual inspection of the data. The HMM detector did not identify any new glitches in pulsars aside from PSR J1413–6141. The high rate of glitches in this pulsar meant a phase connected solution was unable to be obtained without the preliminary glitch properties returned by the HMM detector. Once we had obtained a preliminary solution with confidence in the pulse numbering, we then applied the Bayesian pulsar timing package TEMPONEST (Lentati et al., 2014) to construct posterior probability distributions for both the deterministic and stochastic pulsar properties. Depending whether the pulsar had glitched or not, this final step followed either the single or multi-stage process that we describe below.

4.3.2 Measuring glitch properties

For pulsars that were found to have glitched, we performed an initial TEMPONEST analysis to measure the properties of the glitch(es). The simplest glitches can be described by a permanent step-function in the pulsar spin-frequency, $\Delta\nu_p$. Some glitches required more complex modelling that included step changes in spin-down, $\Delta\dot{\nu}_p$, and one or more exponential recoveries of the spin-frequency towards its pre-glitch value ($\Delta\nu_d$) over time (τ_d) are required to fully describe their phenomenology. The combined, initial step-changes in ν and $\dot{\nu}$ for glitches with recoveries are simply $\Delta\nu_g = \Delta\nu_p + \Delta\nu_d$ and $\Delta\dot{\nu}_g = \Delta\dot{\nu}_p - \Delta\nu_g Q / \tau_d$, where $Q = \Delta\nu_d / \Delta\nu_g$ is the fractional amount by which a glitch has recovered. Since glitches have spin-frequency rise-times of only a few seconds (e.g. Ashton et al., 2019a), few have been detected in the midst of an observation. As a result, there is often some level of ambiguity in both exact time a glitch occurred (t_g) and the precise number of times the neutron star has rotated between the last pre-glitch and first post-glitch observations. This issue can be overcome by adding an unphysical jump in the pulsar rotation phase ($\Delta\phi_g$) to the glitch model, thereby ensuring a phase connected solution is

Table 4.1 Prior ranges on intrinsic, stochastic and glitch parameters. Δ_{param} is the uncertainty returned by TEMPO2, T is length of each pulsar’s data set. The value of x is between 10^3 - 10^5 depending on the pulsar.

Parameter	Prior type (range)
RAJ, DecJ, ν , $\dot{\nu}$, $\ddot{\nu}$ ($^{\circ}$, $^{\circ}$, Hz, s^{-2} , s^{-3})	Uniform ($\pm x \times \Delta_{\text{param}}$)
Proper motion (mas yr^{-1})	Uniform (-1000 , 1000)
EFAC	Uniform (-1 , 2)
EQUAD (s)	Log-uniform (-10 , 1)
Red noise amplitude ($\text{yr}^{3/2}$)	Log-uniform (-15 , -3)
Red noise spectral index	Uniform (0 , 20)
Low frequency cutoff (Hz)	Log-uniform (-1 , 0)
Sinusoid amplitude (s)	Log-uniform (-10 , 0)
Sinusoid phase (rad)	Uniform (0 , 2π)
Log-sinusoid frequency (Hz)	Log-uniform ($1/T_{\text{span}}$, $100/T_{\text{span}}$)
Glitch phase jump (rotations)	Uniform (-5 , 5)
Glitch permanent change in ν (Hz)	Log-uniform (-11 , -4)
Glitch change in $\dot{\nu}$ (Hz^{-2})	Uniform (-10^{-18} , -10^{-11})
Glitch decaying change in ν (Hz)	Log-uniform (-11 , -4)
Glitch recovery timescale (days)	Uniform (1 , 2000)

maintained across the glitch epoch. Adding all of these components together, we obtained the standard model for pulsar rotational phase following a glitch

$$\begin{aligned} \phi_g(t) = & \Delta\phi_g + \Delta\nu_p(t - t_g) + \frac{1}{2}\Delta\dot{\nu}_p(t - t_g)^2 \\ & - \left(\sum_{i=0}^k \Delta\nu_{d,i}\tau_{d,i} \left[1 - e^{-(t-t_g)/\tau_{d,i}} \right] \right), \end{aligned} \quad (4.4)$$

where k is the total number of exponential recoveries. As the exponential is a non-linear function, fitting for it with TEMPO2 requires some level of a priori knowledge of its value. Previous works usually worked around this issue by measuring τ_d from either (or via a combination of) local fits to ν and $\dot{\nu}$ over time or iterative TEMPO2 fits to minimise the χ^2 of the post-fit residuals once a good initial estimate is obtained (e.g. Yu et al., 2013). This non-linearity is less of an issue as we explicitly restricted the prior range of allowable values for τ_d in our parameter estimation. Hence we were able to include the recovery timescale as a free parameter when fitting the ToAs.

To avoid potential biases in our measurements, we modelled both the pulsar glitch and timing noise parameters simultaneously (van Haasteren & Levin, 2013). The red noise

power spectrum is modelled in TEMPONEST as a simple power law of the form

$$P(f) = \frac{A_r^2}{12\pi^2} \left(\frac{f}{\text{yr}^{-1}} \right)^{-\beta_r}, \quad (4.5)$$

in which A_r and β_r are the red noise amplitude (in units of $\text{yr}^{3/2}$) and spectral index and f is the frequency of the power spectrum. Excess scatter in the timing residuals relative to their formal uncertainties due to pulse-to-pulse shape variations and radiometer noise was accounted for by modifying the uncertainties of each ToA as

$$\sigma_{\text{ToA},i}^2 = \sigma_Q^2 + F\sigma_i^2, \quad (4.6)$$

where, σ_Q (the parameter EQUAD in TEMPONEST) accounts for any additional time-independent white noise processes, σ_i is the original uncertainty on the i -th ToA and F (EFAC) is a free-parameter that describes unaccounted instrumental distortions. To reduce the computation time when fitting the glitch parameters, we treated the astrometric and rotational properties of each pulsar as a set of nuisance parameters that are analytically marginalised over. We used the prior ranges listed in Table 4.1 throughout this work. For pulsars that possessed large inter-glitch slopes in $\dot{\nu}$, we had to include a $\ddot{\nu}$ -term when conducting the glitch parameter estimation. Failing to do so resulted in both the $\dot{\nu}$ gradient as well as any $\Delta\dot{\nu}_p$ -terms in the glitch model being partially absorbed by the power-law red noise model. A quick comparison of the Bayesian evidences obtained for pulsars where both timing models were fit revealed an overwhelming preference for the model that included step-changes in the spin-down rate and a $\ddot{\nu}$ term. As an example, for PSR J1420–6048 we obtained a natural log-Bayes factor of 186 in favour of the $\ddot{\nu}$ -inclusive model.

An example outcome of our methodology is shown in Figure 4.1, where the permanent and decaying properties of a glitch in PSR J1524–5625 were well recovered.

4.3.3 Pulsar properties and extended noise models

After obtaining timing models with solutions for the glitches, we then moved on to refining the deterministic properties of each pulsar and searching for additional, previously unmeasured pulsar second spin-period derivatives and proper motion. We generally did not attempt re-fitting glitches in the archival data sets when searching for these new deterministic pulsar properties. Instead, we opted to marginalise over the glitch parameters as many of these glitches had been well characterised in previous works. Our standard approach to measuring the intrinsic, non-glitch properties of each pulsar in our sample included fitting for the pulsar position, rotational and stochastic properties. For pulsars

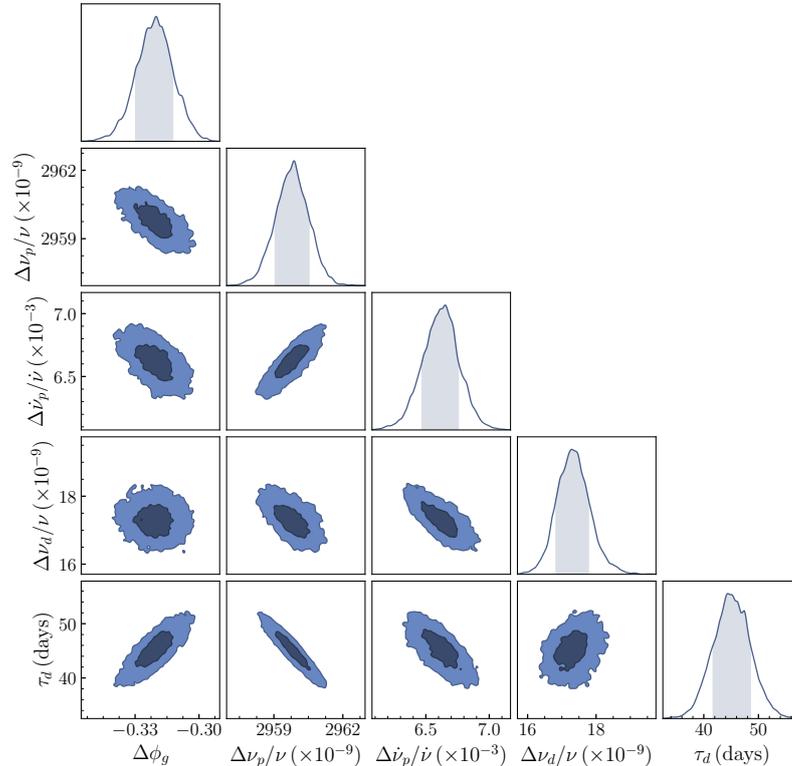


Figure 4.1 Example corner plot showing the one- and two-dimensional posterior distributions for the MJD 55745 glitch in PSR J1524–5625. Contours in the two-dimensional posteriors indicate the 68% and 95% confidence intervals. Shading in the one-dimensional posteriors cover the 68% confidence region.

that have glitched, we used the post-fit ephemeris where the glitch properties were set to be analytically marginalised over during the refitting process. Aside from measuring the standard positional, spin/spin-down and red noise parameters (termed the PL model), we also searched for evidence of additional, unmodelled pulsar properties by iteratively adding them to our timing models. This included fitting for a braking index via including a $\ddot{\nu}$ term (the PL+F2 model) and searching for pulsar proper-motions (PL+PM model) with component values ranging between $\pm 1000 \text{ mas yr}^{-1}$. Following the methodology of Parthasarathy et al. (2019), we also compared various extensions of the standard timing noise model implemented in TEMPO_{NEST}. This included low-frequency components with fluctuation cycles that are longer than our data sets (LFC model), and a variant of the red noise model in Equation 4.5 that incorporates a spectral break in the red power-spectrum (BPL model) as

$$P_r(f) = \frac{A_r^2}{12\pi^2} \frac{(f_c/\text{yr}^{-1})^{-\beta_r}}{[1 + (f/f_c)^{-\beta_r/2}]^2}, \quad (4.7)$$

where f_c is the turnover frequency. Lastly, we checked for the presence of periodic signals that could be induced by the presence of an unmodelled binary companion or quasi-periodic spin-down state switching (e.g. Lyne et al., 2010) by including a sinusoid to the timing model (SIN model). In cases where evidence for more than one process was favoured over the simplest PL model, we conducted more complex joint fits for multiple additional processes. As an example, we might fit for $\ddot{\nu}$, proper-motion and low-frequency components in addition to the other pulsar properties simultaneously – i.e. a PL+F2+PM+LFC model.

We used Bayesian model selection to assess which model best describes the data, specifically through comparing the ratio of Bayesian evidences between competing models, often referred to as the Bayes factor (\mathcal{B}). In this work, we used a relatively conservative Bayes factor threshold of $\ln(\mathcal{B}) = 3$. Under the Kass & Raftery (1995) interpretation of the Bayes factor, $\ln(\mathcal{B}) > 3$ indicates a significant preference for one model over the other. If $\ln(\mathcal{B}) < 3$, then we consider there to be insufficient evidence to distinguish one model from the other.

4.4 Glitches and timing results

We identified and characterised 124 glitches in 52 of the 74 pulsars in our sample. Our measurements of their properties are presented in Table 4.2. This includes 37 previously published glitches, references for which are given in the final column of the table and a further 13 glitches that are publicly listed in the Jodrell Bank glitch catalogue¹ (Espinoza et al., 2011c) but have not been published elsewhere. For the remaining 74 glitches we report their properties for the first time. With the addition of this work, there are now more than 600 glitches known across the pulsar population.

The glitch with the largest amplitude in our sample was that of PSR J1650–4502 on MJD 57780, where $\Delta\nu_g/\nu = (12767_{-1}^{+2}) \times 10^{-9}$. In terms of fractional glitch size, this event is the fourth-largest detected in a rotation-powered pulsar and seventh largest overall when including glitches in magnetars. The recovered step change in $\Delta\dot{\nu}_g/\dot{\nu} = (290_{-60}^{+80}) \times 10^{-3}$ and small fractional glitch recovery of $Q = 0.0061(8)$, are typical of glitches with similar fractional sizes (e.g. Yu et al., 2013). PSR J1650–4502 does appear to be somewhat of an outlier among pulsars that have exhibited very large fractional glitch sizes. It has a \dot{P} that is almost an order of magnitude smaller than the cluster of pulsars that exhibit similarly large amplitude glitches (see panel c of figure 14 in Yu et al. 2013), yet possesses a similar rotation period. A search for flux density and polarization variations following the glitch will be the subject of future work. In contrast, the smallest fractional glitch

¹jb.man.ac.uk/pulsar/glitches/gTable.html

Table 4.2 Median recovered glitch parameters and associated 68% credible intervals (indicated by uncertainties with a + or -). Glitch epochs correspond either to previous reported values, or were computed via the method detailed in Section 3.3 of Espinoza et al. (2011c). Values in parentheses represent the 1- σ uncertainties on the last digit. Glitches with multiple recoveries have additional recovery timescales and fractions listed in rows underneath the main glitch parameters. Note the results listed for PSRs J1341–6220 and J1740–3015 were obtained via TEMPO2 least-squares fitting, where the uncertainties represent to 1- σ uncertainties returned by TEMPO2. See Section 4.4.2 for details regarding glitches 14 and 15 in PSR J1341–6220. References are given to where a glitch was first reported. Those with a \dagger are listed in the Jodrell Bank glitch catalogue but not published elsewhere.

PSRJ	Gl. no.	t_g (MJD)	$\Delta\nu_g/\nu$ ($\times 10^{-9}$)	$\Delta\dot{\nu}_g/\dot{\nu}$ ($\times 10^{-3}$)	$\Delta\dot{\nu}_p/\dot{\nu}$ ($\times 10^{-3}$)	τ_d (days)	Q	Ref
J0631+1036	1	58352.14(4)	120 ± 2	$0.8^{+0.7}_{-0.5}$	$0.8^{+0.7}_{-0.5}$	–	–	This work \dagger
J0729–1448	1	54697(3)	6646^{+13}_{-9}	37^{+25}_{-11}	$11.8^{+0.3}_{-0.2}$	40^{+36}_{-12}	0.006(2)	Weltevrede et al. (2010)
J0742–2822	1	55020.469(4)	100.9 ± 0.3	$0.2^{+0.3}_{-0.2}$	$0.2^{+0.3}_{-0.2}$	–	–	Espinoza et al. (2011c)
J0742–2822	2	56727.7(0.1)	2.6 ± 0.2	$\lesssim 0.2$	$\lesssim 0.2$	–	–	Basu et al. (2020)
J0835–4510	1	55408.802	1902.4 ± 0.5	7 ± 1	6.99 ± 0.09	13 ± 2	0.00548(8)	Buchner (2010)
J0835–4510	2	56555.871	3057 ± 2	4.6 ± 0.3	4.69 ± 0.06	148 ± 8	0.0270(4)	Buchner (2013)
J0835–4510	3	57734.484991(29)	1439.8 ± 0.5	11.2 ± 0.3	6.30 ± 0.03	5.8 ± 0.1	0.00546(8)	Palfreyman et al. (2018)
J0901–4624	1	57179(6)	0.9 ± 0.1	$0.02^{+0.04}_{-0.01}$	$0.02^{+0.04}_{-0.01}$	–	–	This work
J0908–4913	1	58765.06(5)	22.2 ± 0.2	$\lesssim 0.5$	$\lesssim 0.5$	–	–	Lower et al. (2019)
J0940–5428	1	55346(8)	$1573.9^{+1.1}_{-0.8}$	11 ± 2	4.9 ± 0.2	49^{+16}_{-10}	0.0068(8)	This work
J0940–5428	2	58322(16)	$1100.5^{+0.6}_{-0.5}$	$4.0^{+0.2}_{-0.1}$	$4.0^{+0.2}_{-0.1}$	–	–	This work
J1015–5719	1	56695(6)	3232.3 ± 0.6	11 ± 2	3.7 ± 0.2	103^{+19}_{-17}	0.0078(7)	This work
J1016–5857	1	55030(9)	$1919.8^{+1.1}_{-0.9}$	6 ± 1	3.4 ± 0.2	64^{+37}_{-25}	0.005(1)	Yu et al. (2013)
J1016–5857	2	55595(10)	$1464.4^{+1.1}_{-0.9}$	4^{+5}_{-1}	$2.3^{+0.3}_{-0.2}$	88^{+337}_{-50}	0.005(4)	This work
J1016–5857	3	56975(8)	$6.21^{+2.35}_{-1.18}$	–	–	–	–	This work
J1019–5749	1	55595(10)	1.33 ± 0.4	$0.12^{+0.22}_{-0.09}$	$0.12^{+0.22}_{-0.09}$	–	–	This work
J1019–5749	2	55981(10)	377.8 ± 0.4	0.51 ± 0.3	0.51 ± 0.3	–	–	This work
J1028–5819	1	57881(14)	$2296.5^{+0.5}_{-0.4}$	35^{+1}_{-2}	3.5 ± 0.3	54.4^{+2}_{-3}	0.0114(2)	This work
J1048–5832	1	54495(4)	3044.1 ± 0.9	$5.2^{+0.5}_{-0.4}$	$5.2^{+0.5}_{-0.4}$	–	–	This work
J1048–5832	2	56756(4)	2963^{+4}_{-2}	9^{+9}_{-4}	4.0 ± 0.5	33^{+52}_{-17}	0.004(2)	This work
J1052–5954	1	54493.695(1)	6778 ± 1	70^{+7}_{-7}	$6.5^{+0.8}_{-0.7}$	64 ± 6	0.0057(3)	Weltevrede et al. (2010)
J1055–6028	1	57035(10)	$105.7^{+0.6}_{-0.5}$	$1.3^{+0.3}_{-0.2}$	$1.3^{+0.3}_{-0.2}$	–	–	This work
J1105–6107	1	54711(21)	35^{+1}_{-2}	24 ± 4	24 ± 4	–	–	Weltevrede et al. (2010)
J1105–6107	2	55300(16)	949^{+2}_{-1}	17 ± 4	17 ± 4	–	–	Yu et al. (2013)
J1112–6103	1	55288(7)	1793 ± 1	6 ± 2	$3.61^{+1.1}_{-0.8}$	313^{+237}_{-175}	0.014(11)	This work
J1112–6103	2	57922(6)	1283 ± 1	$4.9^{+0.5}_{-0.4}$	$4.9^{+0.5}_{-0.4}$	–	–	This work
J1248–6344	1	56043(5)	1.7 ± 0.2	$\lesssim 0.13$	$\lesssim 0.13$	–	–	This work
J1301–6305	1	55124(10)	4169^{+3}_{-2}	$5.8^{+0.7}_{-0.5}$	$5.8^{+0.7}_{-0.5}$	–	–	This work
J1301–6305	2	57718(6)	658 ± 3	6.0 ± 0.5	6.0 ± 0.5	–	–	This work
J1320–5359	1	56534(10)	10.5 ± 0.1	$0.2^{+0.2}_{-0.1}$	$0.2^{+0.2}_{-0.1}$	–	–	This work
J1320–5359	2	56737(14)	246.8 ± 0.1	$0.08^{+0.14}_{-0.06}$	$0.08^{+0.14}_{-0.06}$	–	–	This work
J1341–6220	1	54468(18)	313 ± 1	0.65 ± 0.08	0.65 ± 0.08	–	–	Weltevrede et al. (2010)
J1341–6220	2	54871(11)	307.2 ± 0.6	1.43 ± 0.06	1.43 ± 0.06	–	–	Weltevrede et al. (2010)
J1341–6220	3	55042(16)	1528 ± 4	63 ± 8	0.36 ± 0.09	9 ± 1	0.042(3)	Yu et al. (2013)
J1341–6220	4	55484(11)	2.4 ± 0.3	0.49 ± 0.01	0.49 ± 0.01	–	–	This work
J1341–6220	5	55835(7)	329.0 ± 0.4	0.98 ± 0.02	0.98 ± 0.02	–	–	This work
J1341–6220	6	56098(12)	151.2 ± 0.4	0.22 ± 0.04	0.22 ± 0.04	–	–	This work
J1341–6220	7	56386(5)	96 ± 1	$\lesssim 0.6$	$\lesssim 0.6$	–	–	This work
J1341–6220	8	56479(9)	37 ± 2	$\lesssim 0.6$	$\lesssim 0.6$	–	–	This work
J1341–6220	9	56602(12)	1709.7 ± 0.3	1.06 ± 0.1	1.06 ± 0.1	–	–	This work
J1341–6220	10	57357(9)	28.7 ± 0.7	0.25 ± 0.09	0.25 ± 0.09	–	–	This work
J1341–6220	11	57492(1)	21.7 ± 0.1	0.8 ± 0.1	0.8 ± 0.1	–	–	This work

Table 4.2 (Continued)

PSRJ	Gl. no.	t_g (MJD)	$\Delta\nu_g/\nu$ ($\times 10^{-9}$)	$\Delta\dot{\nu}_g/\dot{\nu}$ ($\times 10^{-3}$)	$\Delta\dot{\nu}_p/\dot{\nu}$ ($\times 10^{-3}$)	τ_d (days)s	Q	Ref
J1341–6220	12	57647(13)	706 ± 1	0.77 ± 0.09	0.77 ± 0.09	–	–	This work
J1341–6220	13	57880(14)	60 ± 3	$\lesssim 0.6$	$\lesssim 0.2$	86 ± 47	0.102(3)	This work
J1341–6220	14	58178(15)	*	*	*	–	–	This work
J1341–6220	15	58214(4)	*	*	*	–	–	This work
J1357–6429	1	54803(17)	2332^{+4}_{-3}	13 ± 1	13 ± 1	–	–	Weltevrede et al. (2010)
J1357–6429	2	55576(10)	4860^{+3}_{-2}	$14.7^{+0.7}_{-0.8}$	$14.7^{+0.7}_{-0.8}$	–	–	This work
J1357–6429	3	57795(22)	2250 ± 11	7 ± 2	7 ± 2	–	–	This work
J1357–6429	4	58148(15)	1930^{+5}_{-4}	$\lesssim 0.9$	$\lesssim 0.9$	–	–	This work
J1406–6121	1	56193(10)	$143.5^{+1.6}_{-0.8}$	2^{+4}_{-2}	$0.4^{+0.3}_{-0.2}$	107^{+206}_{-62}	0.03(2)	This work
J1410–6132	1	54652(19)	263 ± 2	$\lesssim 0.23$	$\lesssim 0.23$	–	–	Weltevrede et al. (2010)
J1413–6141	1	54303(1)	2412 ± 3	$\lesssim 0.6$	$\lesssim 0.6$	–	–	Yu et al. (2013)
J1413–6141	2	55744(7)	235 ± 2	$\lesssim 0.7$	$\lesssim 0.7$	–	–	This work
J1413–6141	3	56147(12)	200^{+3}_{-2}	$\lesssim 0.3$	$\lesssim 0.3$	–	–	This work
J1413–6141	4	56567(5)	371^{+2}_{-1}	$\gtrsim 0.8$	$\gtrsim 0.8$	–	–	This work
J1413–6141	5	56975(8)	30 ± 2	$\gtrsim 0.4$	$\gtrsim 0.4$	–	–	This work
J1413–6141	6	57236(4)	266 ± 2	$\gtrsim 0.9$	$\gtrsim 0.9$	–	–	This work
J1413–6141	7	57509(10)	356 ± 2	$\gtrsim 0.4$	$\gtrsim 0.4$	–	–	This work
J1413–6141	8	57838(6)	2137 ± 2	$\gtrsim 0.5$	$\gtrsim 0.5$	–	–	This work
J1420–6048	1	54652(9)	$927.6^{+0.7}_{-0.6}$	6 ± 1	4.19 ± 0.4	45^{+20}_{-16}	0.007(4)	Weltevrede et al. (2010)
J1420–6048	2	55410(9)	$1352.8^{+0.5}_{-0.4}$	5.4 ± 0.2	5.4 ± 0.2	–	–	Yu et al. (2013)
J1420–6048	3	56267(6)	1954.2 ± 0.3	5.7 ± 0.2	5.7 ± 0.2	–	–	This work
J1420–6048	4	57210(8)	1210^{+2}_{-1}	9^{+7}_{-4}	$3.5^{+0.4}_{-0.3}$	19^{+25}_{-12}	0.009(2)	This work
J1452–6036	1	55055.22(4)	28.95 ± 0.03	$2.5^{+1.0}_{-0.9}$	$\lesssim 0.7$	2340^{+822}_{-672}	0.12(5)	Yu et al. (2013)
J1452–6036	2	57115(6)	0.13 ± 0.03	$\lesssim 0.2$	$\lesssim 0.2$	–	–	This work
J1452–6036	3	58600.292(3)	270.61 ± 0.03	1.2 ± 0.3	1.2 ± 0.3	–	–	Lower et al. (2020c)
J1524–5625	1	55745(7)	$2977.0^{+0.7}_{-0.5}$	$15.5^{+0.9}_{-0.7}$	$6.6^{+0.2}_{-0.1}$	45^{+4}_{-3}	0.0058(2)	This work
J1614–5048	1	55734(2)	4098 ± 3	3 ± 1	3 ± 1	–	–	This work
J1614–5048	2	56443(11)	5949^{+9}_{-7}	$\lesssim 1$	$\lesssim 1$	–	–	This work
J1617–5055	1	54747(7)	334 ± 3	9^{+11}_{-2}	0.48^{+5}_{-1}	227^{+262}_{-38}	0.975(6)	This work
J1617–5055	2	55164(9)	11 ± 2	0.8 ± 0.6	0.8 ± 0.6	–	–	This work
J1617–5055	3	55316(6)	68 ± 2	$2.2^{+0.6}_{-0.5}$	$2.2^{+0.6}_{-0.5}$	–	–	This work
J1617–5055	4	55873(11)	55 ± 2	1.1 ± 0.6	1.1 ± 0.6	–	–	This work
J1617–5055	5	56267(6)	2068 ± 2	$13.2^{+0.6}_{-0.7}$	$13.2^{+0.6}_{-0.7}$	–	–	This work
J1644–4559	1	56600(14)	$717.4^{+0.3}_{-0.2}$	$0.4^{+0.5}_{-0.3}$	$0.4^{+0.5}_{-0.3}$	–	–	This work
J1646–4346	1	55288(7)	8591^{+6}_{-5}	16^{+9}_{-5}	$8.1^{+1.2}_{-0.9}$	126^{+137}_{-62}	0.005(2)	This work
J1650–4502	1	57778(8)	12767^{+2}_{-1}	290^{+80}_{-60}	27^{+7}_{-6}	82^{+21}_{-16}	0.0061(8)	This work
J1702–4128	1	57719(6)	3090 ± 1	10^{+2}_{-1}	4.77 ± 0.09	88^{+32}_{-19}	0.0040(4)	This work
J1702–4310	1	57510(10)	3129^{+4}_{-1}	5^{+2}_{-1}	$3.4^{+0.1}_{-0.2}$	50^{+72}_{-36}	0.002(1)	This work
J1705–3950	1	58236(14)	9331^{+6}_{-4}	61^{+11}_{-10}	$5.9^{+0.3}_{-0.2}$	55^{+10}_{-9}	0.0053(4)	This work
J1709–4429	1	54711(22)	2752.5 ± 0.2	$13.8^{+0.9}_{-1.0}$	7.4 ± 0.09	55^{+6}_{-7}	0.010(1)	Weltevrede et al. (2010)
J1709–4429	2	56354(13)	2951.9 ± 0.6	8 ± 1	$4.2^{+0.4}_{-0.3}$	54^{+11}_{-9}	0.006(1)	This work
J1709–4429	3	58178(6)	$2432.8^{+0.7}_{-0.6}$	8.5 ± 0.9	$4.6^{+0.4}_{-0.3}$	49 ± 8	0.0061(9)	Lower et al. (2018)
J1718–3825	1	54911(2)	2.2 ± 0.2	$\lesssim 0.08$	$\lesssim 0.08$	–	–	Yu et al. (2013)
J1718–3825	2	57950(7)	7.1 ± 0.1	$\lesssim 0.07$	$\lesssim 0.07$	–	–	This work
J1730–3350	1	55926(6)	$2250.7^{+1.0}_{-0.9}$	7^{+3}_{-2}	5 ± 2	151.01^{+199}_{-81}	0.007(1)	This work [†]
J1731–4744	1	55735.18(14)	52.7 ± 0.4	3 ± 1	$0.51^{+0.5}_{-0.4}$	151^{+35}_{-57}	0.10(5)	Shternin et al. (2019)
J1731–4744	2	56239.86(77)	11.0 ± 0.3	$\lesssim 0.4$	$\lesssim 0.4$	–	–	Shternin et al. (2019)
J1731–4744	3	56975(8)	6.4 ± 0.3	$\lesssim 0.2$	$\lesssim 0.2$	–	–	This work

Table 4.2 (Continued)

PSRJ	Gl. no.	t_g (MJD)	$\Delta\nu_g/\nu$ ($\times 10^{-9}$)	$\Delta\dot{\nu}_g/\dot{\nu}$ ($\times 10^{-3}$)	$\Delta\dot{\nu}_p/\dot{\nu}$ ($\times 10^{-3}$)	τ_d (days)	Q	Ref
J1731–4744	4	57978.17(2)	3145^{+2}_{-4}	$1.2^{+0.4}_{-0.5}$	$1.2^{+0.4}_{-0.5}$	–	–	Jankowski et al. (2017)
J1734–3333	1	56351(12)	$86.7^{+7.9}_{-7.1}$	$0.3^{+0.3}_{-0.2}$	$0.3^{+0.3}_{-0.2}$	–	–	This work [†]
J1737–3137	1	54348(4)	1341.8 ± 0.6	$3.0^{+2.0}_{-0.8}$	$1.7^{+0.2}_{-0.3}$	152^{+224}_{-75}	0.004(1)	Espinoza et al. (2011c)
J1737–3137	2	57147(8)	8.1 ± 0.4	$\lesssim 0.1$	$\lesssim 0.1$	–	–	This work
J1737–3137	3	58207(29)	$4494.1^{+0.9}_{-1.4}$	1.20 ± 0.08	1.20 ± 0.08	–	–	This work
J1737–3137	4	58838(24)	15 ± 2	$\lesssim 0.1$	$\lesssim 0.1$	–	–	This work
J1740–3015	1	54450.19(1)	45 ± 2	$\lesssim 2$	$\lesssim 0.7$	54 ± 9	0.15(1)	Weltevrede et al. (2010)
J1740–3015	2	54695.19(2)	2 ± 1	$\lesssim 0.5$	$\lesssim 0.5$	–	–	Yuan et al. (2010)
J1740–3015	3	54810.9(1)	4 ± 1	$\lesssim 0.2$	$\lesssim 0.2$	–	–	Espinoza et al. (2011c)
J1740–3015	4	54928.6(1)	3 ± 0.7	$\lesssim 0.1$	$\lesssim 0.1$	–	–	Espinoza et al. (2011c)
J1740–3015	5	55220(14)	2659 ± 4	1.2 ± 0.2	0.7 ± 0.1	258 ± 31	0.008(1)	Yu et al. (2013)
J1740–3015	6	55936.2(1)	16 ± 6	1.1 ± 0.3	1.1 ± 0.3	–	–	This work [†]
J1740–3015	7	57499.371(4)	228 ± 3	2.2 ± 0.5	0.97 ± 0.03	89 ± 13	0.057(9)	Jankowski et al. (2016)
J1740–3015	8	58232.4(4)	835.5 ± 0.4	0.26 ± 0.04	0.26 ± 0.04	–	–	Basu et al. (2020)
J1757–2421	1	55702(6)	$7800.5^{+2.1}_{-0.4}$	66^{+31}_{-29}	3.1 ± 0.2	15^{+10}_{-9}	0.0003(2)	Yuan et al. (2017)
						92^{+23}_{-12}	0.0014(1)	
						618^{+215}_{-108}	0.0021(3)	
J1801–2304	1	55371.1(2)	3.0 ± 0.5	$\lesssim 0.18$	$\lesssim 0.18$	–	–	Yu et al. (2013)
J1801–2304	2	55851.7(5)	$1.6^{+0.6}_{-0.7}$	$\lesssim 0.16$	$\lesssim 0.16$	–	–	This work [†]
J1801–2304	3	56158.360(2)	513.6 ± 0.5	$\lesssim 0.25$	$\lesssim 0.25$	–	–	This work [†]
J1801–2304	4	57586.4(1)	96.2 ± 0.5	$\lesssim 0.21$	$\lesssim 0.21$	–	–	This work [†]
J1801–2451	1	54661(2)	3083.7 ± 0.7	6.5 ± 0.5	6.5 ± 0.5	–	–	Yu et al. (2013)
J1801–2451	2	56943(7)	2423.5 ± 0.9	$5.9^{+0.5}_{-0.4}$	$5.9^{+0.5}_{-0.4}$	–	–	This work
J1803–2137	1	55775(2)	$4785.9^{+1.2}_{-0.9}$	17 ± 1	7.3 ± 0.3	40^{+5}_{-4}	0.0071(5)	This work [†]
J1825–0935	1	53734.6(1)	6 ± 2	$1.2^{+1.9}_{-0.9}$	$1.2^{+1.9}_{-0.9}$	–	–	Espinoza et al. (2011c)
J1825–0935	2	54115.78(4)	117^{+13}_{-11}	1^{+2}_{-1}	1^{+2}_{-1}	–	–	Yuan et al. (2010)
J1826–1334	1	56534(10)	129.6 ± 0.2	$1.27^{+0.11}_{-0.08}$	$1.27^{+0.11}_{-0.08}$	–	–	This work
J1826–1334	2	56690(9)	2421.2 ± 0.3	$4.5^{+0.7}_{-0.5}$	$4.0^{+0.5}_{-0.4}$	164^{+162}_{-87}	0.002(1)	This work
J1837–0604	1	55873(11)	1376 ± 1	8 ± 3	$1.5^{+1.2}_{-0.8}$	328^{+125}_{-100}	0.06(2)	This work [†]
J1837–0604	2	56503(4)	$16.7^{+0.8}_{-0.7}$	$\lesssim 0.7$	$\lesssim 0.7$	–	–	This work [†]
J1841–0345	1	58455(7)	$112.3^{+0.8}_{-0.7}$	$56.3^{+0.83}_{-6.1}$	$56.3^{+0.83}_{-6.1}$	–	–	This work
J1841–0524	1	54503(21)	1032.8 ± 0.5	$\lesssim 2$	$\lesssim 0.1$	488^{+162}_{-150}	0.023(9)	Weltevrede et al. (2010)
J1841–0524	2	55524(19)	$806.2^{+0.6}_{-0.5}$	$\lesssim 2$	$\lesssim 0.1$	288^{+112}_{-87}	0.018(5)	This work [†]
J1841–0524	3	56567(2)	23 ± 1	$\lesssim 0.5$	$\lesssim 0.5$	–	–	This work [†]
J1847–0402	1	55509(1)	0.47 ± 0.03	$0.02^{+0.03}_{-0.01}$	$0.02^{+0.03}_{-0.01}$	–	–	This work [†]
J1847–0402	2	58244(21)	0.16 ± 0.04	$0.01^{+0.02}_{-0.01}$	$0.01^{+0.02}_{-0.01}$	–	–	This work

we observed was found in PSR J1452–6036 on MJD 57115. With an amplitude of only $\Delta\nu_g/\nu = 0.13 \pm 0.03 \times 10^{-9}$, this glitch is among the 16 smallest glitches ever recorded.

4.4.1 Minimum detectable glitch size and sample completeness

A comparison of how our sample of 124 glitches fit in with those previously published is shown in Figure 4.2, where our measurements of $\Delta\nu_g$ and $\Delta\dot{\nu}_g$ span nearly the complete spectrum of reported values in v1.64 of the ATNF glitch catalogue (Manchester et al.,

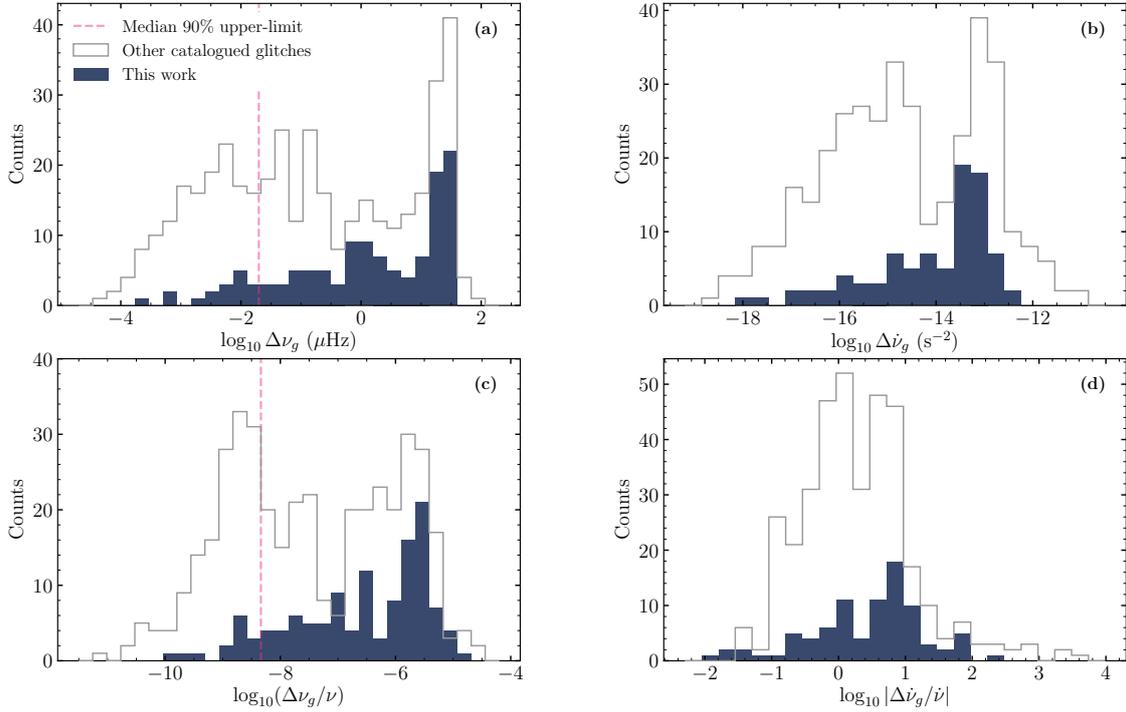


Figure 4.2 Histograms of $\Delta\nu_g$ (a), $\Delta\dot{\nu}_g$ (b), $\Delta\nu_g/\nu$ (c) and $\Delta\dot{\nu}_g/\dot{\nu}$ (d) for the 124 glitches in our sample (‘This work’; filled dark blue) and the larger ATNF glitch catalogue after removing those listed in Table 4.2 (open grey). The dashed magenta lines in panels (a) and (c) indicate the median upper-limit on the glitch size after averaging across our pulsar sample (see text for details).

2005)². The general lack of glitches below $\Delta\nu_g \lesssim 10^{-2} \mu\text{Hz}$, or $\Delta\nu_g/\nu \lesssim 10^{-9}$ in the bottom panel, is likely the result of a selection effect where smaller glitches become increasingly difficult to detect because of limitations in pulsar observing cadence. Indeed, the relatively sparse observation cadence of the P574 programme (approximately one observation every month) makes it difficult to differentiate the smallest glitches ($\Delta\nu_g/\nu \lesssim 1 \times 10^{-9}$) from stochastic variations in pulse phase due to timing noise. As a result, three previously catalogued small glitches were not picked up by the HMM detector in PSR J1740–3015 along with the glitch in PSR J1825–0935. Note, the non-detection of the glitch in PSR J1825–0935 may be a result of it occurring within a 651 d gap in the Parkes timing observations. However, targeted parameter estimation with TEMPO₂ was able to recover all but one of the small glitches in PSR J1740–3015, and the missed glitch in PSR J1825–0935.

A natural question to ask is: for a given data set, what is the smallest glitch which

²<https://www.atnf.csiro.au/research/pulsar/psrcat/glitchTbl.html>

can be reliably detected? We tackled this question empirically using synthetic data sets. For reasons of practicality we could only approach the question this way using the HMM detector – even for a single data set many hundreds of injections are necessary to give a reasonable estimate of the performance of the detector, and so an automated approach to glitch detection is required. The essential figure of merit is a 90% frequentist upper limit on the size of undetected glitches in each data set, denoted $\Delta\nu_g^{90\%}$ (see Appendix B for details). We calculated a $\Delta\nu_g^{90\%}/\nu$ for the full timing baseline of nearly every pulsar considered in the sample, as well as a separate $\Delta\nu_g^{90\%}/\nu$ across each stretch of inter-glitch data. A complete list of the $\Delta\nu_g^{90\%}/\nu$ values can be found in Table B1.

There is considerable variation in $\Delta\nu_g^{90\%}/\nu$ across the sample with values between 1.4×10^{-10} and 7.7×10^{-8} . This is attributable to differences in the amount of timing noise in the individual pulsars. Variations may also arise due to differences in observing cadence, particularly the presence of long gaps in the data. As a result, we obtain mean and median upper limits of $\Delta\nu_g^{90\%}/\nu = 8.1 \times 10^{-9}$ and $\Delta\nu_g^{90\%}/\nu = 4.6 \times 10^{-9}$ respectively across the entire sample.

4.4.2 Vela pulsar timing and pulsars with more than five glitches

PSR J0835–4510 (B0833–45, ‘Vela’): The Vela pulsar was the first pulsar observed to glitch (Radhakrishnan & Manchester, 1969) and has been observed for over five decades. Our attempts to fit the three large glitches experienced by Vela within our timing data via TEMPONEST did not converge, which may be due to the numerical precision issue that also affected the Shannon et al. (2016) analysis of 21 yrs of Vela pulsar timing. While Shannon et al. (2016) worked around this issue by implementing long-double precision in a bespoke version of TEMPONEST, implementing a similar correction was not practicable for our analysis. Instead, we performed generalised least-square fitting to the Vela pulsar ToAs to measure its glitch properties using TEMPO2. We added additional glitch parameters to the model in an iterative fashion, where parameters were only kept in the final model if the weighted root-mean-square of the residuals was lower than before they were included. In Table 4.2 we report improved measurements of $\Delta\nu_g$ and $\Delta\dot{\nu}_g$ for the two glitches on MJDs 55408 and 56555 over the previously reported values in Buchner (2010) and Buchner (2013). Only single recovery processes were measured for all three glitches, with decay timescales of 13 ± 2 d, 149 ± 8 d and 5.9 ± 0.1 d respectively. Our timing cadence was insufficient to resolve the short-term 0.96 ± 0.17 d recovery reported by Sarkissian et al. (2017). The additional long-term recovery processes that Shannon et al. (2016) found evidence for could not be constrained, largely due to our inability to model the glitches and

timing noise simultaneously via the least-squares fitting of TEMPO2. Improved modelling of the Vela glitches is left for future work.

PSR J1341–6220 (B1338–62): This pulsar has a low characteristic age ($\tau_c = 12.1$ kyr) and is potentially associated with the supernova remnant G308.8–0.1 (Kaspi et al., 1992). It is the most prolific glitching pulsar in our sample, with 24 previously published glitches found between MJD 47989 and 55484. In addition to re-analysing three glitches previously found by Weltevrede et al. (2010) and Yu et al. (2013), we report the discovery of a further 12 new glitches that occurred between MJD 55484 and 58214. The prevalence of glitches in this pulsar precluded the use of our standard TEMPONEST-based approach to inferring the glitch properties, as the large dimensionality of the timing model resulted in even highly parallelised TEMPONEST runs failing to converge. Instead, we computed generalised least-squares fits to the ToAs using TEMPO2. As a result, our recovered glitch parameters will be slightly contaminated by unaccounted red noise. The 14th and 15th glitches are separated by only a single observing epoch, hence we could only obtain a joint measurement of their properties, where $\Delta\nu_g/\nu = (155.4 \pm 0.5) \times 10^{-9}$ and $\Delta\dot{\nu}_g/\dot{\nu} = (1.47 \pm 0.03) \times 10^{-3}$. We are confident there are indeed two closely spaced glitches as opposed to a single glitch as there is a clear detection of the pulsar during the intervening epoch, and our attempts to fit for only a single glitch assuming the epoch of glitch 14 failed to whiten the residuals. Glitch 9 is the largest to be reported in this pulsar to date, and only the fifth found to have a fractional amplitude greater than $\Delta\nu_g/\nu = 10^{-6}$, while glitch 4 has the smallest amplitude of any found in this pulsar.

PSR J1413–6141: We report seven new glitches in this pulsar with a variety of amplitudes ranging between $\Delta\nu_g/\nu = (30 \pm 2) \times 10^{-9}$ to $(2137 \pm 2) \times 10^{-9}$. These glitches are smaller than the largest reported glitch in this pulsar on MJD 54303 (Yu et al., 2013), where our re-analysis recovered an amplitude of $\Delta\nu_g/\nu = (2412 \pm 3) \times 10^{-9}$. The fifth glitch in our sample is the smallest to have been found in this pulsar to date, smaller than the $\Delta\nu_g/\nu = (39 \pm 4) \times 10^{-9}$ glitch on MJD 51290 (Yu et al., 2013). None of these glitches appear to have induced detectable step-changes in $\dot{\nu}$, nor show evidence for exponential recoveries in ν or $\dot{\nu}$.

PSR J1740–3015 (B1737–30): As for PSR J1341–6220, the large number of glitches necessitated using the generalised least-squares method implemented in TEMPO2 to measure their properties. Of the eight previously published glitches, only seven were recovered with values of $\Delta\nu_g$ that are inconsistent with zero. We only obtained an upper-limit of $\Delta\nu_g/\nu \lesssim 0.9 \times 10^{-9}$ for the small glitch reported by Jankowski et al. (2016) as occurring on MJD 57346. This glitch was not picked up by the HMM glitch detection

algorithm. It was also not recovered in a re-analysis of the UTMOST timing data by Lower et al. (2020c), suggesting this event may have been a misidentified variation in spin-phase caused by timing noise. The only unpublished glitch within our timing baseline for PSR J1740–3015 occurred on MJD 55936 with a moderate amplitude of $\Delta\nu_g/\nu = (16 \pm 6) \times 10^{-9}$, $\Delta\dot{\nu}_g/\dot{\nu} = (1.1 \pm 0.3) \times 10^{-9}$ and no apparent exponential recovery.

Statistics of glitches in PSRs J1341–6220 and J1413–6141

Recent developments in modelling the statistics of pulsar glitches have focused on microphysics-agnostic meta-models. Two such meta-models, where stress is accumulated either as a state-dependent Poisson process (SDP; Carlin & Melatos, 2019) or as a Brownian process (BSA; Carlin & Melatos, 2020), are predicted to show similar auto- and cross-correlations between their glitch amplitudes and wait-times. However, under the fast-driven SDP, pulsars are expected to have glitch amplitude and wait-time distributions with the same overall shape (Carlin & Melatos, 2019), whereas the amplitudes and wait-times can be drawn from differing distributions under the BSA.

Using the sample of 34 well constrained glitches, we find PSR J1341–6220 exhibits Spearman rank correlation coefficients of $\rho_{s+} = 0.67 \pm 0.14$ (p-value, 0.33 ± 0.01) and $\rho_{s-} = -0.18 \pm 0.18$ (p-value, $(2_{-3}^{+1}) \times 10^{-5}$) for the cross-correlations between the glitch size and the forward (referring to the previous glitch) and backward (referring to the next glitch) waiting times respectively. Note, the p-value for the backward cross-correlation is an artefact of the significant scatter and corresponding lack of correlation. The moderate forward cross-correlation is consistent with previously reported values (Melatos et al., 2018; Fuentes et al., 2019; Carlin & Melatos, 2020). The marginal auto-correlations for the glitch wait time, $\rho_{s, \Delta T_g} = -0.0048$ (p-value 0.98), and size, $\rho_{s, \Delta\nu} = -0.26$ (p-value 0.15), are consistent with the BSA process.

The 14 glitches in PSR J1413–6141 show a strong forward cross-correlation between the wait time and sizes of its glitches, with $\rho_{s+} = 0.82 \pm 0.17$ (p-value, $0.29_{-0.04}^{+0.02}$), and a weak backwards anti-correlation of $\rho_{s-} = -0.31 \pm 0.29$ (p-value, $(6_3^{+2}) \times 10^{-4}$). Similar to PSR J1341–6220, the auto-correlations of PSR J1413–6141 are relatively weak, with Spearman coefficients for the wait time and amplitude of $\rho_{s, \Delta T_g} = -0.0070$ (p-value 0.98) and $\rho_{s, \Delta\nu} = -0.27$ (p-value 0.36) respectively.

4.4.3 Recoveries

Of the 124 pulsar glitches listed in Table 4.2 only 41 were found to have at least one exponential recovery where the recovery parameters are well constrained. Figure 4.3 shows

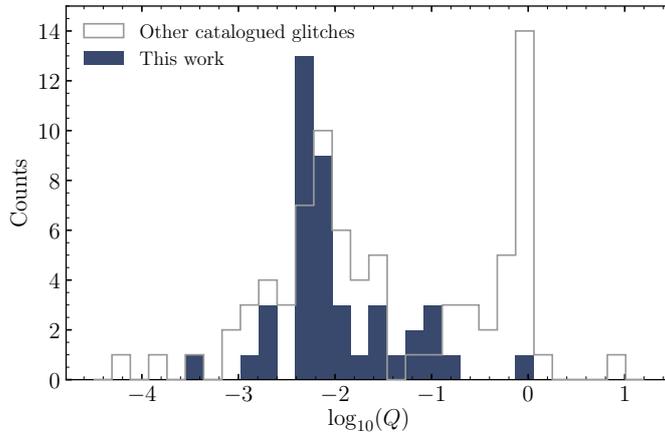


Figure 4.3 Same as Figure 4.2 but for the distribution of fractional glitch recoveries (Q).

the fractional recovery of the spin-frequency following a glitch. Our sample possesses only a single peak at $Q \sim 0.0069$, unlike the bi-modal distribution seen in the overall glitch sample, but we note that the second peak at high Q -values largely comprises of recoveries measured in the Crab pulsar (PSR J0534+2200; Lyne et al. 2015), which is not observed by our programme. The only glitch from our set that has high Q is the MJD 54762 glitch in PSR J1617–5055.

For the large glitch in PSR J1757–2421 on MJD 55702, our parameter estimation recovered the ~ 15 - and ~ 92 -day glitch recovery timescales found by Yuan et al. (2017) in addition to a third, long-term exponential recovery process with $\tau_{d,3} \sim 618$ days. This three-component recovery model is strongly preferred over one containing only two-components, with $\ln(\mathcal{B}) = 10.3$. The presence of this additional recovery process would explain the stronger post-glitch timing noise they measured. Each of the three exponential recoveries exhibited sequentially larger recovery fraction, a phenomenon that is seen in the short- and long-term recoveries of the Vela pulsar (e.g. Shannon et al., 2016). Our total fractional recovery of $Q = 0.0038(2)$ is consistent with the value of $Q = 0.0035(9)$ found by Yuan et al. (2017), however their measurement only included the first two recoveries. This could be due the difference in methods that were applied for measuring the recovery terms, as we fit for all of the glitch parameters simultaneously as opposed to incrementally adding and fitting for additional parameters.

4.4.4 Glitch statistics

An empirical relation for determining the glitch rate for a given pulsar was developed by Fuentes et al. (2017). They divided glitches into ‘large’ and ‘small’ with a separation at

$\Delta\nu_g = 10\mu\text{Hz}$. As the completeness to small glitches is not well established and the size distribution of large glitches is relatively tight (see Figure 4.2), they were able to derive an approximate wait time, T_g , between large glitches via

$$T_g = \frac{1}{420 \text{ Hz}^{-1} \dot{\nu}}. \quad (4.8)$$

We can test this on our sample of pulsars. Combining the data sets from Parthasarathy et al. (2019) and this paper, we have observed 159 pulsars over a 10 year span. From Equation 4.8, assuming Poisson statistics for the glitch waiting times, our expectation is for 25 pulsars to undergo a total of 45 glitches for an average realisation. This compares well with the actual value of 28 pulsars having experienced 44 large glitches.

When we rank the pulsars by expected glitch activity, PSR J1513-5908 is highest, with the expectation of a glitch every 1.1 yr. In fact, this pulsar has not had a single glitch (large or small) in more than 35 yr of observing (Parthasarathy et al., 2020). Apart from this singular anomaly, 18 of the next 20 pulsars in rank order have all had large glitches. We also note that of the eight pulsars with multiple large glitches all have values of $T_g < 12$ yr. In our sample of pulsars with a large glitch, PSR J1757-2421 has the longest wait time of 320 yr, or a glitch probability in 10 years of only 0.0028 according to Equation 4.8. It is clear that Equation 4.8 is a simplification of the underlying processes which determine glitch activity (see e.g. Melatos et al. 2018), and assumes that large and small glitches arise from different processes. Nevertheless, it does a reasonable job at predicting wait-times for pulsars with characteristic ages below 10^5 yr.

4.4.5 New and updated pulsar parameters

In addition to fitting for the rotational properties of our pulsars, we also generated posterior distributions with TEMPO_{NEST} for their astrometric properties. Tables containing the recovered astrometric and rotational properties of each pulsar can be found in Appendix C.

As a result, we measured the proper-motions of 11 pulsars, summarized in Table 4.3. Here the proper-motion in RA and DEC are given by $\mu_\alpha \equiv \dot{\alpha} \cos(\delta)$ and $\mu_\delta \equiv \dot{\delta}$, which when combined gives the total proper-motion $\mu_T = \sqrt{\mu_\alpha^2 + \mu_\delta^2}$. The 2-D (or 1-D in the case of a proper-motion in only a single coordinate) transverse velocity is calculated as

$$V_T = 4.74 \text{ km s}^{-1} \left(\frac{\mu_T}{\text{mas yr}^{-1}} \right) \left(\frac{D}{\text{kpc}} \right), \quad (4.9)$$

where D is the distance to the pulsar. For most pulsars we used the median dispersion

Table 4.3 Pulsars with significant proper motions. Those with proper motions measured for the first time are highlighted in bold. Unless otherwise specified, the distances reported here are inferred from pulsar dispersion measures. *Distance from parallax measurement. †Average distance from electron density models and HI-observations.

PSRJ	μ_α (mas yr ⁻¹)	μ_δ (mas yr ⁻¹)	μ_T (mas yr ⁻¹)	D (kpc)	V_T (km s ⁻¹)
J0659+1414	49 ± 13	78 ⁺⁷⁵ ₋₆₈	96 ⁺⁷⁰ ₋₄₅	3.47*	132 ⁺⁹⁶ ₋₆₂
J0908–4913	–37 ± 9	31 ± 10	47 ± 9	3.0†	674 ± 127
J1003–4747	–12 ± 2	21 ± 2	24 ± 3	0.37	42 ± 3
J1057–5226	49 ± 4	–6 ± 5	50 ± 4	0.09	21 ± 2
J1320–5359	13 ± 2	52 ± 2	54 ± 2	2.2	563 ⁺²⁴ ₋₂₁
J1359–6038	–4 ± 4	10 ± 5	12 ⁺⁵ ₋₄	5.0	278 ⁺¹⁰⁹ ₋₁₀₀
J1452–6036	–5 ± 3	–5 ± 3	7 ± 3	6.1	204 ⁺⁹⁶ ₋₈₈
J1709–4429	17 ± 3	11 ± 10	17 ± 4	2.6	210 ± 45
J1731–4744	60 ⁺¹¹ ₋₉	–178 ⁺²⁴ ₋₂₂	183 ± 23	0.7	607 ± 75
J1826–1334	32 ⁺⁸ ₋₉	–	32 ⁺⁸ ₋₉	3.6	549 ⁺¹⁴⁴ ₋₁₆₂

measure distance returned by the TC93 (Taylor & Cordes, 1993), NE2001 (Cordes & Lazio, 2002) and YMW16 (Yao et al., 2017) Galactic free-electron density models. Where available, we used distances inferred from the parallax measurements (e.g. Deller et al., 2019).

Of the four new measurement of proper motion, those of PSRs J1359–6038 and J1452–6036 are well constrained in μ_α and μ_δ , with both pulsars possessing V_T that are consistent with the bulk pulsar population (see Hobbs et al., 2005) at their nominal distance. For PSR J0908–4913 (B0906–49), the posterior distributions for μ_α and μ_δ are shown in Figure 4.4. The pulsar is moving in a north-westerly direction, which is consistent with the implied direction of motion from radio imaging of its bow-shock nebula (Gaensler et al., 1998). The implications of this proper motion are discussed further in a companion paper (Johnston & Lower, 2021). Finally for PSR J1003–4747 there is a large disparity in its distance inferred via the NE2001 and YMW16 electron-density models, and hence its implied V_T . The NE2001 model gives a distance of 2.9 kpc with a corresponding $V_T = 335 \pm 27 \text{ km s}^{-1}$, whereas the YMW16 distance of only 0.37 kpc gives $V_T = 42 \pm 3 \text{ km}^{-1}$. A velocity measurement via scintillation or a direct distance measurement via parallax would solve this issue.

Our measurements for pulsars with previously reported proper motions are largely consistent with the published values. This includes PSR J1709–4429, where our radio

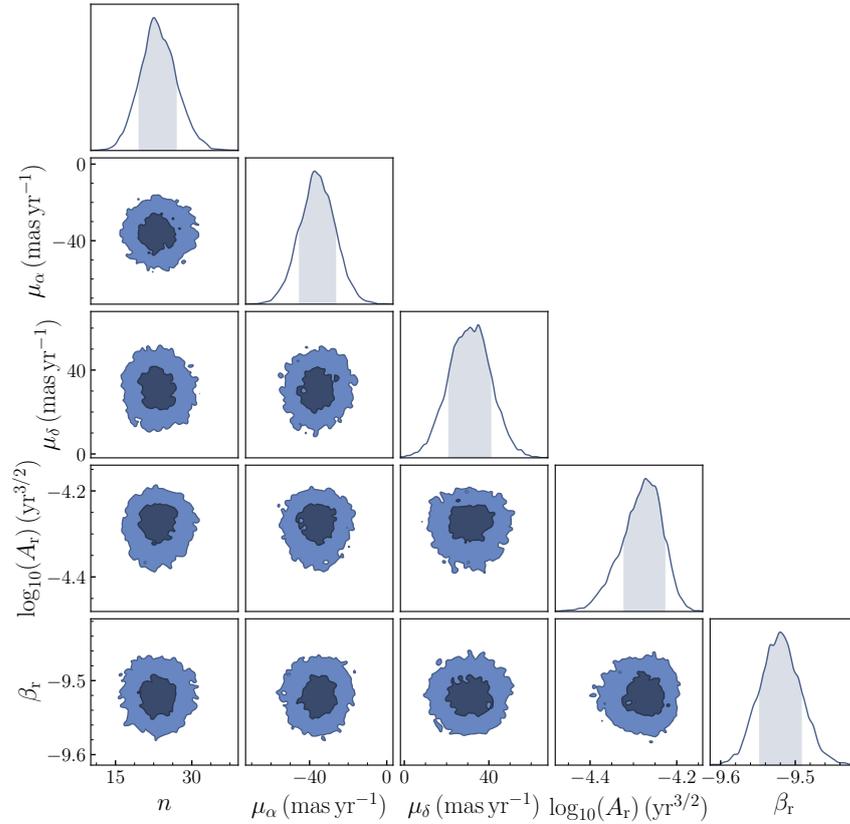


Figure 4.4 One- and two-dimensional posterior distributions for the braking index, proper motion and red noise parameters of PSR J0908–4913.

measurements are consistent with the values of $\mu_\alpha = 13 \pm 2$ and $\mu_\delta = 1 \pm 2$ obtained from X-ray imaging (de Vries et al., 2021). The proper-motion of PSR J1731–4744 (B1727–47) is consistent with the values of $\mu_\alpha = 73 \pm 15 \text{ mas yr}^{-1}$ and $\mu_\delta = -132 \pm 14 \text{ mas yr}^{-1}$ inferred by Shternin et al. (2019), adding further credence to the claimed association with the supernova remnant RCW 114.

Finally, PSR J1617–5055 has a small characteristic age (8 kyr) and high rotational kinetic energy loss rate ($\dot{E} = 1.6 \times 10^{37} \text{ ergs s}^{-1}$) but no pulsed gamma-ray emission has yet been detected from the pulsar. One suggestion for this absence was the lack of a coherent timing solution due to the large amount of timing noise (Abdo et al., 2013). However, even our phase coherent timing solution was unable to recover pulsed gamma-ray emission in *Fermi* data covering MJD 54220–56708 (D. A. Smith 2020, private communication). As observations of PSR J1617–5055 were discontinued after MJD 56708 our timing model is unable to be used for folding additional *Fermi* photons, as any subsequent glitches are unaccounted for.

4.5 Braking indices and long-term evolution

Of the 74 pulsars in our sample, we found 34 were best described by a timing model that included a power-law red noise process and a $\ddot{\nu}$ -term. The resulting Bayes factors in favour of the model containing a $\ddot{\nu}$ -term over one where $\ddot{\nu}$ is fixed at zero, along with the inferred n from Equation 4.2 and observing timespan are listed in Table 4.4. A graphical comparison of the braking indices and associated 68 per cent confidence intervals from both this work and Parthasarathy et al. (2020) is depicted in Figure 4.5.

Parthasarathy et al. (2020) reported a weak correlation between n and τ_c , with a Spearman coefficient of $\rho_s = 0.34 \pm 0.01$ amongst their sample of 19 predominately non-glitching pulsars. We find there is no correlation among our sample of pulsars ($\rho_s = 0.09$, p-value 0.61) nor in a joint analysis of both samples ($\rho_s = 0.15$, p-value 0.28). Hence the weak correlation of Parthasarathy et al. (2020) may have simply been a result of both a selection effect – i.e. avoidance of generally younger, glitching pulsars – and an artefact of their relatively small sample size.

Quantitatively, we assessed whether the two sets of braking index measurements satisfy the null hypothesis, i.e. that they were drawn from the same underlying distribution, by performing a two-sample Kolmogorov-Smirnov (KS) test on the cumulative distributions of $\log_{10}(n)$. Omitting PSR J0857–4424, which has a large n perhaps resulting from an unmodelled binary companion in a wide orbit (Parthasarathy et al., 2020), we obtain $\mathcal{D}_{KS} = 0.16$ and a p-value of 0.15. Here, the KS-statistic that is lower than the critical value of $\mathcal{D}_{0.05} = 0.45$ and p-value that is > 0.05 indicates the null hypothesis cannot be rejected, and both distributions are indistinguishable from one another at the 0.05 level. If the braking indices of both glitching and non-glitching pulsars are indeed drawn from a common distribution, then the same underlying mechanism may be exerting the torque.

4.5.1 Observed versus predicted evolution in P - \dot{P} space

Measurements of ν , $\dot{\nu}$ and n can be used to make predictions on how the rotation of a pulsar may evolve over long timescales. Re-writing Equation 4.1 in terms of P and \dot{P} , the motion of a pulsar in the P - \dot{P} diagram can be described as

$$\dot{P} = \kappa P^{2-n}, \quad (4.10)$$

where pulsars with $n < 2$ will move towards larger values of both P and \dot{P} over time, while pulsars with increasingly larger values of n evolve more rapidly towards smaller \dot{P} values. A key assumption here is that κ remains constant over time. If this is not

Table 4.4 Braking indices for pulsars with $\ln(\mathcal{B}_{F2=0}^{F2 \neq 0}) > 3$. Those highlighted with a ‘★’ have been observed to glitch. Pulsars above the line have either never been seen to glitch, or do not have large values of $\Delta\dot{\nu}_g/\dot{\nu}$ ($\lesssim 1 \times 10^{-3}$) associated with their glitches. For pulsars below the line, the listed values of n are representative of their average inter-glitch braking, not their long-term evolution.

PSRJ	$\ln(\mathcal{B}_{F2=0}^{F2 \neq 0})$	Braking index (n)	T (d)
J0659+1414★	21.0	$12.8_{-0.2}^{+0.3}$	3964
J0855–4644	16.8	$7.8_{-0.2}^{+0.3}$	4249
J0901–4624★	122.2	$13.4_{-0.9}^{+1.0}$	7620
J0908–4913★	11.8	23_{-3}^{+4}	10112
J1320–5359★	14.2	111_{-14}^{+16}	7933
J1410–6132★	8.2	22 ± 3	4116
J1718–3825★	57.8	$48.7_{-0.8}^{+0.7}$	7593
J1726–3530	30.5	19 ± 2	6028
J1734–3333★	5.4	1.2 ± 0.2	5986
J1841–0425★	14.5	189 ± 18	4202
J0835–4510★	227.2	44 ± 2	4209
J0940–5428★	49.9	30 ± 1	8353
J1015–5719★	3.6	16_{-1}^{+2}	4249
J1016–5857★	49.7	23_{-1}^{+2}	7170
J1028–5819★	14.9	58_{-10}^{+12}	3906
J1048–5832★	65.0	32 ± 2	10560
J1112–6103★	10.3	42 ± 6	7620
J1301–6305★	88.0	25.2 ± 0.7	7504
J1357–6429★	190.5	38.2 ± 0.5	6978
J1420–6048★	255.5	$47.8_{-1.0}^{+0.9}$	7175
J1524–5625★	73.6	43.2 ± 0.7	4249
J1614–5048★	56.1	$14.3_{-0.6}^{+0.3}$	10140
J1617–5055★	19.8	33_{-9}^{+8}	2488
J1646–4346★	17.2	29 ± 2	10557
J1702–4128★	8.8	12.7 ± 0.6	4250
J1702–4310★	80.7	$13.8_{-0.4}^{+0.5}$	7248
J1709–4429★	187.7	$35.2_{-0.5}^{+0.7}$	10561
J1730–3350★	20.7	$20.7_{-1.5}^{+0.7}$	7783
J1731–4744★	29.6	54_{-4}^{+2}	9427
J1737–3137★	5.5	15 ± 1	4752
J1801–2451★	345.3	$38.2_{-0.9}^{+1.0}$	9573
J1803–2137★	102.6	$32.0_{-0.7}^{+0.8}$	4250
J1826–1334★	33.2	32 ± 1	4184

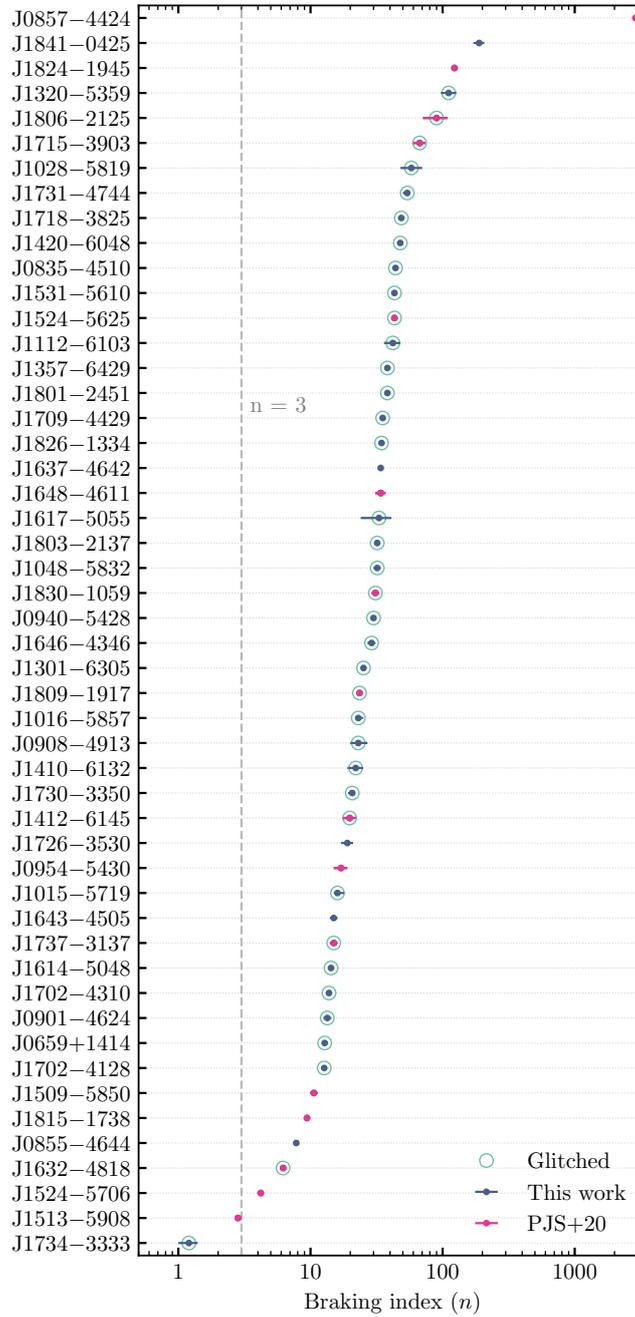


Figure 4.5 Median recovered braking indices and 68% confidence intervals for the 33 pulsars in Table 4.4 (dark blue) and Parthasarathy et al. (2020) (PJS+20; magenta). Glitching pulsars are highlighted by green circles.

the case, then the measured n from Equation 4.2 would represent the ensemble effect of both the intrinsic braking process and a process that causes κ to vary with time.

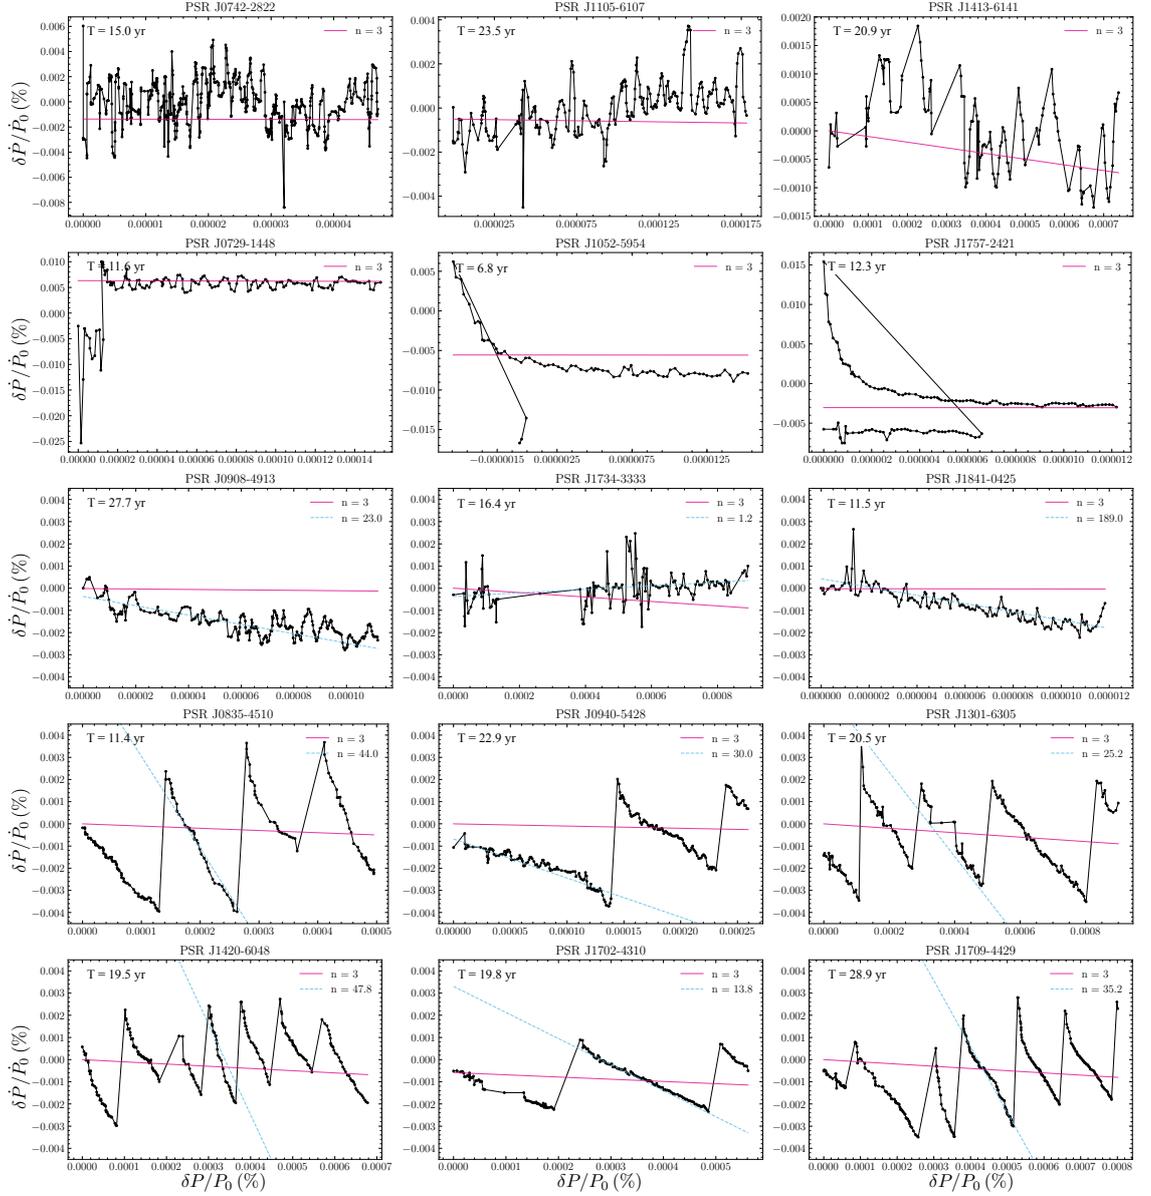


Figure 4.6 Observed long-term, fractional evolution of P and \dot{P} for a sample of 15 pulsars. The organisation of the panels follows the logic described in Section 4.5.1. Expected evolutionary paths due to dipole radiation ($n = 3$) and (where possible) the measured braking indices are indicated by the solid magenta and dashed blue lines respectively. Note these lines do not represent fits to the data. See text for further details.

Additionally, discontinuities and recovery effects induced by glitches also cause pulsars to deviate from expected evolutionary paths. This kind of behaviour has been clearly seen in young pulsars such as the Crab pulsar, where the jumps in $\dot{\nu}$ associated with its glitches result in a smaller braking index of $n = 2.342 \pm 0.001$ compared to the glitch-corrected

value of $n = 2.519 \pm 0.002$ (Lyne et al., 2015).

In Figure 4.6 we show the observed fractional evolution of P and \dot{P} for 15 pulsars that represent the various evolutionary pathways that are visible throughout our overall sample. Each of these fractional P - \dot{P} diagrams were produced by performing stride-fits to ν and $\dot{\nu}$ over discrete windows containing at least five ToAs, and then converting the resulting values to spin-period ($P = 1/\nu$) and period derivative ($\dot{P} = -\dot{\nu}/\nu^2$). The timing models used for performing the local fits did not account for the glitches seen in these pulsars. For pulsars with a measured $\ddot{\nu}$ component, we overlaid both the expected evolutionary lines for $n = 3$ and their observed braking indices. Pulsars where $\ddot{\nu}$ could not be distinguished have only the $n = 3$ evolutionary track overlaid. We separated the pulsars into four broad groups based on the visual appearance of their evolutionary tracks in P - \dot{P} space.

1. Flat: pulsars with a small, undetected n that appear to evolve with roughly constant \dot{P} over our timing baseline (top row of Figure 4.6)
2. Flat-jumps: pulsars with an unresolved n with jumps in \dot{P} associated with large glitches (second row of Figure 4.6).
3. Inclined: pulsars that follow an evolutionary track defined by a constant n over the duration of our timing programme (third row of Figure 4.6).
4. Vela-like: pulsars that exhibit both large, positive inter-glitch braking indices and (quasi-)periodic jumps in \dot{P} due to large glitches (bottom two rows of Figure 4.6).

Pulsars in Table 4.4 that experienced glitches with $\Delta\dot{\nu}_g/\dot{\nu} \gtrsim 10^{-3}$ components belong to the Vela-like group. They evolve with large n between glitches, with the $\Delta\dot{\nu}_g$ component of the glitch serving to ‘reset’ much of their \dot{P} evolution towards the value observed immediately after the previous glitch. Indeed, examination of the panels of Figure 4.6 shows the expected evolutionary tracks corresponding to values of n from Table 4.4 matches the observed inter-glitch behaviour, while the long-term evolution of these pulsars is consistent with a ‘small n ’ process, though the precise value of this long-term n is unclear. This behaviour is in line with previous studies of $\dot{\nu}$ evolution among samples of actively glitching, young pulsars (see figure 1 in Espinoza et al., 2017). Measurements of the n that dominates their long-term evolution would require simultaneous measurements of both the effects of glitches, pulsar rotation and astrometry, and the addition of new $\Delta\ddot{\nu}_p$ parameters to account for their large inter-glitch braking indices. Such a high-dimensional problem is incompatible with our current approach to modelling the timing of these pulsars, and is therefore left for future works.

The remaining pulsars from Table 4.4 fall into the inclined group of pulsars. Unlike the Vela-like pulsars, these pulsars have yet to undergo a large glitch within our data span, and they therefore follow the evolutionary path given by their measured n over at least the timescales covered by the Parkes timing programme. The pulsars that have undergone small glitches (such as PSR J0908–4913 and PSR J1734–3333) continue along their pre-glitch paths without interruption over the decades of available timing data. Note that the y-axis scale of the inclined and Vela-like pulsars in Figure 4.6 is identical, and we see that for the glitching pulsars, the maximum deviation of \dot{P} from the $n = 3$ line is approximately 0.004%. The inclined pulsars, however, have not yet had time to reach this critical point of their inter-glitch evolution and so we predict that they will undergo large glitches similar to the Vela-like pulsars at some point in the (near) future. We surmise that the pulsars with high n in Parthasarathy et al. (2020) fall into the same category. While there does appear to be some level of variation in $\ddot{\nu}$ between glitches in the Vela-like pulsars shown in Figure 4.6, the precise nature of the relationship between the glitches and the measured braking indices is unclear. This aspect is further explored in Section 4.5.2.

Pulsars not listed in Table 4.4 can be categorised as possessing evolutionary paths that fall into either the flat or flat-jump categories. In general, these pulsars do not show strong evidence for a $\ddot{\nu}$ term in their timing model over a model with red noise alone, although this can be explained given their relatively small $\dot{\nu}$ values in relation to other pulsars in our sample. Hence, longer observing time spans may be required to resolve the $\ddot{\nu}$ component of their rotational evolution. The younger pulsars (where $\tau_c \lesssim 10^4$ yr) generally undergo small ($\Delta\nu_g/\nu < 10^{-6}$) glitches with negligible $\Delta\dot{\nu}_g$ components (with a handful of exceptions), while the older flat-jump pulsars have singular, extremely large amplitude glitches with a significant $\Delta\dot{\nu}_g$. For instance, PSRs J0729–1448 and J1413–6141, both of which are presented in Figure 4.6, have characteristic ages of 35 and 14 kyr respectively. Despite experiencing several large glitches (see Table 4.2 and Yu et al. 2013), PSR J1413–6141 did not exhibit any resolved changes in $\dot{\nu}$. When combined with a lack of a distinguishable $\ddot{\nu}$, it appears to follow a flat evolutionary path like many of the older, non-glitching pulsars. In contrast, PSR J0729–1448 clearly falls into the flat-jumps category thanks to a large amplitude glitch with a strong $\Delta\dot{\nu}$ component. This large glitch is similar to the giant glitches found in PSRs J1052–5954, J1650–4502 and J1757–2421, which have characteristic ages of 143, 376 and 285 kyr respectively and underwent the three largest glitches listed in Table 4.2. The observed P - \dot{P} evolution of both PSRs J1052–5954 and J1757–2421 are almost entirely dominated by exponential recoveries that occurred

following their glitches, whereas the glitch in PSR J0729–1448 did not show any evidence of recovery.

There is also the question of whether the flat and flat-jumps pulsars are related to the inclined and Vela-like pulsars. Given the comparatively low values of $\dot{\nu}$ associated with most of the flat pulsars, their inferred glitch wait times from Equation 4.8 are significantly longer than their current timing baselines – much like the inclined pulsars. Physically, this can be ascribed to the low spin-down rates of these pulsars requiring longer periods of time to build-up a sufficiently large stress within the neutron star for a glitch to be triggered. Under this scenario, our detection of the giant glitches in the flat-jump pulsars can be attributed to both the large size and hundreds of years’ worth of pulsar timing accumulated by our sample. There are of course a number of obvious exceptions, namely PSR J1413–6141 that we discussed earlier in this section. The potential implications of these pulsars are further expanded upon in Section 4.5.3 below.

The braking index of PSR J1734–3333

PSR J1734–3333 is a pulsar with a high-magnetic field strength that was found by Espinoza et al. (2011b) to possess an unusually small braking index of $n = 0.9 \pm 0.2$. They surmised that this small n may be evidence that the magnetic field strength of this pulsar may be growing over time. Similarly small braking indices can also result from particle outflows (e.g. Michel & Tucker, 1969). The results of our TEMPONEST analysis of this pulsar are shown in Figure 4.7, where the recovered $n = 1.2 \pm 0.2$ is consistent with the Espinoza et al. (2011b) value at the 68 percent confidence interval. PSR J1734–3333 did undergo a relatively small glitch on MJD 56350 with no detectable post-glitch recovery. We tested whether the glitch had any effect on n by conducting a separate TEMPONEST run that excluded the post-glitch ToAs. The resulting posterior distribution for n was almost identical to what we obtained when including the post-glitch data, with a Jensen-Shannon divergence (Lin, 1991)³ of $\lesssim 0.002$ bit. Hence our measurement is unaffected by any low-level glitch recovery effects that were not modelled. As noted in Espinoza et al. (2011b), and is obvious in our Figure 4.6, the pulsar appears to be headed towards the location of the magnetars in $P-\dot{P}$ space. If this braking index were to remain constant over time, then the pulsar would take around (38 ± 9) kyr to obtain similar rotational properties to the bulk magnetar population.

³Identical probability distributions have a Jensen-Shannon divergence of 0 bit (i.e. no information gained between distributions) and maximally divergent distributions have 1 bit.

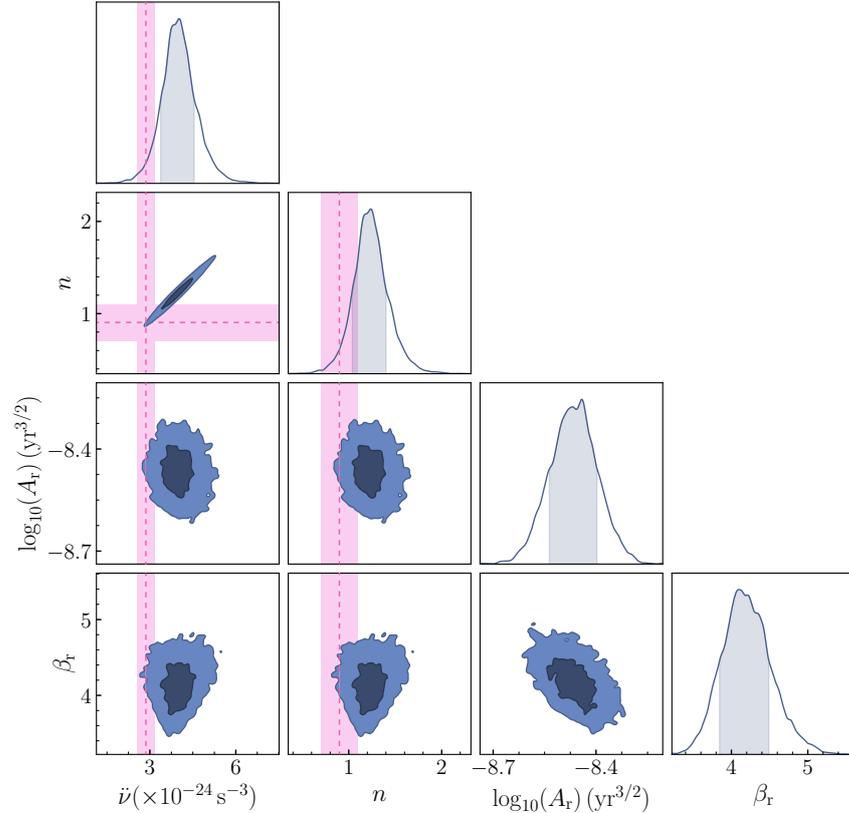


Figure 4.7 One- and two-dimensional posterior distributions for $\dot{\nu}$, braking index and red noise parameters of PSR J1734–3333. Dashed magenta lines and shading indicate the $\dot{\nu}$ and n values and $1-\sigma$ uncertainties from Espinoza et al. (2011c).

4.5.2 Connecting large braking indices to glitches in Vela-like pulsars

In order to understand the relationship between the glitch parameters and our inferred values of n listed in Table 4.4, we investigated how n varied before, after and between consecutive glitches. To measure the separate inter-glitch braking indices (n_{int}), we followed a variation of the technique employed by Yu et al. (2013). The ToAs for the pulsars in Table 4.4 with two or more glitches were separated into multiple sections (pre-first-glitch, inter-glitch and post-final-glitch), while the remaining pulsars that experienced only a single glitch were split into pre- and post-glitch sections. Using TEMPONEST, we fitted for ν , $\dot{\nu}$ and $\ddot{\nu}$ along with at least one set of glitch recovery parameters (ν_d and τ_d) to each section of ToAs. The inter-glitch fits were performed independent of the ‘global’, long-term fits referred to in Section 4.5.2. Hence, modelling of the white and red noise were also performed during each inter-glitch fit. Table 4.5 presents the resulting n_{int} , wait times between glitches (T_g) and both the preceding glitch amplitude and fractional change

in $\dot{\nu}$. Lower limits on T_g are listed for entries corresponding to the pre-first-glitch/post-last-glitch and single glitch cases. As with the inferred braking indices presented in Table 4.4, the values of n_{int} are all positive because they are robust to the presence of the timing noise, which is fitted simultaneously as a power-law process. Note the measurements obtained here are representative of how the pulsars are behaving between glitches, and do not necessarily reflect their long-term rotational evolution (cf. Figure 4.6).

As noted in Section 4.5.1, the pulsars that display large amplitude glitches with a significant $\Delta\dot{\nu}_g$ component all possess large n . There are several ways to interpret this behaviour. The glitches could be viewed as a mechanism that serves to reset much of the rapid downward P - \dot{P} evolution experienced by these pulsars back toward a longer-term, ‘low- n ’ evolutionary track. However, in most microphysical theories, glitches are triggered by stress accumulation in ν (e.g. differential rotation between crust and superfluid) rather than $\dot{\nu}$, so one expects glitches to be triggered by high torque rather than high $n \propto \ddot{\nu}$. On the other hand, n_{int} could stem from a form of post-glitch recovery. Suggested mechanisms include Ekman circulation in a two-component star (van Eysden & Melatos, 2010, 2012), changes in the effective moment of inertia from progressive re-coupling of the crust to the superfluid core (e.g. Smith, 1999; Antonopoulou et al., 2018; Pizzochero et al., 2020; Montoli et al., 2020), unpinning and re-pinning of vortices between crustal pinning sites due to thermal fluctuations (the ‘vortex creep’ model; Alpar et al., 1984b,a, 1993; Alpar & Baykal, 2006; Akbal et al., 2017) and turbulence within an array of vortices pinned in the superfluid (Melatos & Peralta, 2007; van Eysden & Melatos, 2012; Melatos & Link, 2014; Haskell et al., 2020). Although these theoretical models are difficult to falsify (Haskell & Melatos, 2015), and it is unclear whether the type of assumed vortex creep can be supported within physical neutron stars (Link, 2014), they do provide a set of phenomenological behaviour that we can test. In the particular instance of the vortex creep model, $\ddot{\nu}_{\text{int}}$ reflects the gradient of $\dot{\nu}(t)$ as it undergoes a linear recovery. In the lead-up to the next glitch. This can be seen in the phenomenological model of Alpar & Baykal (2006), as equating their equations 11 and 12 returns

$$n_{\text{int}} = 2 \times 10^{-3} \left(\frac{\tau_c}{T_g} \right) \left(\frac{\Delta\dot{\nu}_g/\dot{\nu}}{10^{-3}} \right). \quad (4.11)$$

Substituting both our Equation 4.2 and $\tau_c = 0.5 \nu |\dot{\nu}|^{-1}$ into this relation and re-arranging for $\ddot{\nu}_{\text{int}}$ gives

$$\ddot{\nu}_{\text{int}} = \frac{\Delta\dot{\nu}_g}{T_g}. \quad (4.12)$$

Akbal et al. (2017) assumes a variant of this relation in their modelling of the Vela pulsar,

Table 4.5 Non-zero braking indices measured before, after and between subsequent glitches for the 16 pulsars with one or more large glitches, associated waiting times (or lower-limits) until the next glitch, as well as the amplitudes and fractional changes in spin-down of the preceding glitch.

PSRJ	n_{int}	T_g (days)	$(\Delta\nu_g/\nu)_{\text{prev}}$	$(\Delta\dot{\nu}_g/\dot{\nu})_{\text{prev}}$
J0835–4510	$33.3^{+0.7}_{-0.6}$	> 1130	N/A	N/A
J0835–4510	51.6 ± 0.5	1164	1902.4 ± 0.5	7 ± 1
J0835–4510	52^{+4}_{-3}	1176	3057 ± 2	4.6 ± 0.3
J0835–4510	62 ± 2	> 734	1908.3 ± 0.2	11.2 ± 0.3
J0940–5428	21 ± 1	> 4479	N/A	N/A
J0940–5428	57^{+5}_{-3}	2925	$1573.9^{+1.1}_{-0.8}$	11 ± 2
J1015–5719	16 ± 1	> 2462	N/A	N/A
J1015–5719	87^{+12}_{-6}	> 1761	3232.3 ± 0.6	11 ± 2
J1016–5857	25 ± 1	2491	1622.6 ± 0.3	3.69 ± 0.05
J1016–5857	35 ± 3	1515	$1919.8^{+1.1}_{-0.9}$	6 ± 1
J1016–5857	132^{+29}_{-34}	421	$1464.4^{+1.1}_{-0.9}$	4^{+5}_{-1}
J1048–5832	58 ± 6	1754	2995 ± 7	3.7 ± 0.1
J1048–5832	47 ± 5	1945	771 ± 2	4.62 ± 0.06
J1048–5832	64^{+16}_{-12}	940	1838.4 ± 0.5	3.7 ± 0.3
J1048–5832	37^{+5}_{-6}	817	28.5 ± 0.4	0.19 ± 0.14
J1048–5832	32 ± 3	2260	3044.1 ± 0.9	$5.2^{+0.5}_{-0.4}$
J1320–5359	104^{+13}_{-14}	> 5975	N/A	N/A
J1320–5359	1847^{+347}_{-434}	197	10.5 ± 0.1	$0.2^{+0.2}_{-0.1}$
J1320–5359	151^{+36}_{-26}	> 1703	246.8 ± 0.1	$0.08^{+0.14}_{-0.06}$
J1357–6429	18.1 ± 0.2	2712	2332^{+4}_{-3}	13 ± 1
J1357–6429	73.4 ± 0.9	843	4860^{+3}_{-2}	$14.7^{+0.7}_{-0.8}$
J1357–6429	34.8 ± 0.8	2219	2250 ± 11	7 ± 2
J1420–6048	$49.9^{+0.9}_{-1.0}$	1154	1146.2 ± 0.6	3.83 ± 0.08
J1420–6048	45 ± 2	971	2019 ± 10	6.6 ± 0.8
J1420–6048	60^{+7}_{-6}	947	1270 ± 3	3.9 ± 0.3
J1420–6048	74 ± 1	757	$927.6^{+0.7}_{-0.6}$	6 ± 1
J1420–6048	56 ± 2	838	$1352.8^{+0.5}_{-0.4}$	5.4 ± 0.2
J1420–6048	48 ± 3	944	1954.2 ± 0.3	5.7 ± 0.2
J1420–6048	49 ± 2	> 1241	1210^{+2}_{-1}	9^{+7}_{-4}
J1524–5625	43.0 ± 0.7	> 1509	N/A	N/A
J1524–5625	43 ± 1	> 2711	$2975.9^{+0.7}_{-0.6}$	$15.5^{+0.9}_{-0.7}$
J1617–5055	32^{+10}_{-8}	519	68 ± 2	$2.2^{+0.6}_{-0.5}$

Table 4.5 (Continued)

PSRJ	n_{int}	T_g (days)	$(\Delta\nu_g/\nu)_{\text{prev}}$	$(\Delta\dot{\nu}_g/\dot{\nu})_{\text{prev}}$
J1617–5055	31^{+24}_{-36}	358	55 ± 2	1.1 ± 0.6
J1646–4346	75^{+19}_{-22}	1324	885 ± 3	1.5 ± 0.3
J1702–4310	11^{+2}_{-1}	> 2551	N/A	N/A
J1702–4310	$15.2^{+0.5}_{-0.6}$	3562	4810 ± 27	17 ± 4
J1702–4310	23 ± 2	> 939	3129^{+4}_{-1}	5^{+2}_{-1}
J1709–4429	16 ± 1	> 835	N/A	N/A
J1709–4429	$27.0^{+0.7}_{-0.4}$	2713	2057 ± 2	4.0 ± 0.1
J1709–4429	69 ± 5	1228	1166.7 ± 0.2	6.22 ± 0.03
J1709–4429	41.2 ± 0.6	1977	2872 ± 7	8.0 ± 0.7
J1709–4429	$27.7^{+0.7}_{-0.4}$	1661	2755 ± 1	$13.8^{+0.9}_{-1.0}$
J1709–4429	$36.8^{+0.6}_{-0.5}$	1824	3027^{+7}_{-4}	8 ± 1
J1709–4429	181 ± 31	> 262	$2433.5^{+0.8}_{-0.6}$	8.5 ± 0.9
J1718–3825	48.9 ± 0.4	3004	2.2 ± 0.2	$\lesssim 0.08$
J1718–3825	26 ± 2	> 505	7.1 ± 0.1	$\lesssim 0.7$
J1730–3350	22^{+5}_{-6}	> 1530	N/A	N/A
J1730–3350	$16.4^{+0.3}_{-0.6}$	3871	3202 ± 1	5.9 ± 0.1
J1730–3350	47 ± 4	> 2381	$2250.7^{+1.0}_{-0.9}$	$7+3_{-2}$
J1801–2451	48^{+4}_{-3}	1164	1987.9(3)	4.6 ± 0.1
J1801–2451	50^{+7}_{-8}	1384	1247.4(3)	4.7 ± 0.2
J1801–2451	41 ± 13	921	3755.8(4)	6.8 ± 0.1
J1801–2451	25^{+4}_{-3}	1613	17.4 ± 0.2	1.4 ± 0.1
J1801–2451	26 ± 2	2239	3083.7 ± 0.7	6.5 ± 0.5
J1801–2451	41 ± 3	> 1511	2423.5 ± 0.9	$5.9^{+0.5}_{-0.4}$

though modified to allow for small permanent shifts in $\dot{\nu}$. The turbulent vortex array model of [Haskell et al. \(2020\)](#) also does not attempt to directly relate the internal physics to the observed properties of the neutron star. However, they do predict that $\dot{\nu}_{\text{int}}$ (and by extension n_{int}) should follow a quadratic dependence on T_g if the vortices are pinned within a turbulent region of the star, or a linear dependence for a straight vortex array. [van Eysden & Melatos \(2010\)](#) showed that Ekman pumping in a cylindrical vessel containing a two-component superfluid can result in a non-linear spin-down after applying an impulsive acceleration, i.e. an effective $n \gg 3$ process (see their figure 2). An extension of this model was successful in replicating much of the observed post-glitch behaviour of the Crab and Vela pulsars ([van Eysden & Melatos, 2012](#)), though no explicit formula relating the

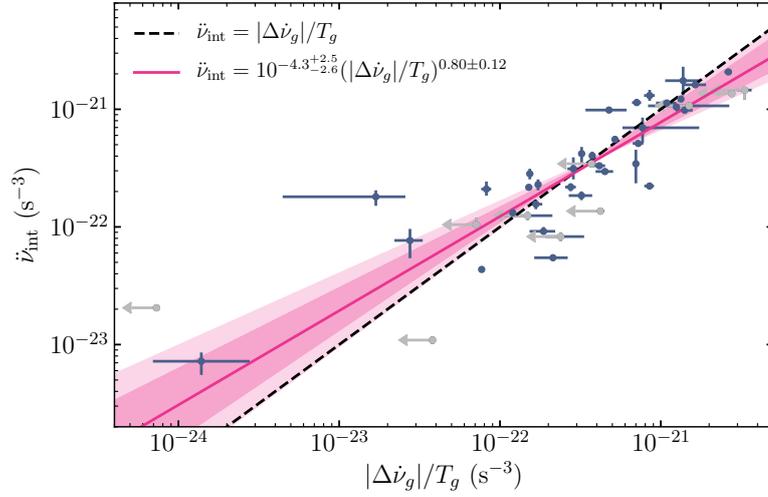


Figure 4.8 Comparison between $\ddot{\nu}_{\text{int}}$ and the $\Delta\dot{\nu}_g$ of the previous glitch divided by the inter-glitch wait time for 53 glitches in 16 pulsars. Black-dashed line indicates an exact one-to-one relationship, while the magenta line and shading represents our median power-law fit and the 68 and 95 percent confidence regions respectively.

model to n_{int} was derived.

On an individual basis, there was no clear correlation between n_{int} and the glitch properties for the pulsars with more than one inter-glitch measurement. However, examining at the sample as a whole revealed a weak anti-correlation between n_{int} (and by extension, $\ddot{\nu}_{\text{int}}$) and T_g , with $\rho_s = -0.65$ (p-value, 2.8×10^{-5}). This correlation weakens with the inclusion of lower-limits on T_g , as well as long-term braking indices from the inclined pulsars, where the total observing span for them is taken as a lower-limit on T_g . Including both the extra terms from Table 4.5 for the inclined pulsars, and employing a bootstrap approach to sample the lower-limits, returned a Spearman coefficient of $\rho_s = -0.37 \pm 0.08$ (p-value, $0.003^{+0.008}_{-0.002}$). Despite the weakness of these anti-correlations, owing to the significant scatter relative to the sample size, the implied inverse dependence of n_{int} on T_g does not match the predictions of Haskell et al. (2020) for straight or turbulent vortices. The most striking evidence for a potential connection between n_{int} and glitches arises from our comparison of the values $\ddot{\nu}_{\text{int}}$ from Table 4.5 with $\Delta\dot{\nu}_g$ divided by T_g in Figure 4.8. It is clear the expected one-to-one relation from the linear recovery process (Equation 4.12) indicated by the dashed line is largely adhered to. Fitting a power-law to only the points with confident measurements of $|\Delta\dot{\nu}_g|/T_g$ using a projected, bivariate Gaussian likelihood

(see equations 26 through 32 of [Hogg et al. 2010](#)), we obtained the relation

$$\ddot{\nu}_{\text{int}} = 10^{-4.3^{+2.5}_{-2.6}} (|\Delta\dot{\nu}_g|/T_g)^{0.80 \pm 0.12}, \quad (4.13)$$

which when plotted in [Figure 4.8](#) is consistent with exact one-to-one relation at the 95 percent confidence interval. We also checked whether there was a stronger forwards or backwards correlation between $\ddot{\nu}_{\text{int}}$ and $\Delta\dot{\nu}_g/\dot{\nu}$ (i.e, comparing to $\Delta\dot{\nu}_g$ from the preceding or following glitch). Using the same bootstrapping technique from earlier, we obtained overlapping Spearman coefficients of $\rho_{s-} = 0.74 \pm 0.04$ for the $\Delta\dot{\nu}_g$ of previous glitch and $\rho_{s+} = 0.82^{+0.03}_{-0.06}$ for the next glitch. The consistency between the two is not surprising given the fractional step-change in spin-down and inter-glitch wait-times are similar among the Vela-like pulsars (see [Figure 4.6](#) and [Table 4.2](#)).

4.5.3 Pulsars with seemingly small braking indices

While there is a clear link between large glitches and large values of n_{int} in the Vela-like pulsars, the same cannot be said for the pulsars that experienced predominately small glitches ($\Delta\nu_g/\nu < 10^{-6}$) and those that have undergone single gigantic glitches ($\Delta\nu_g/\nu \sim 10^{-5}$). The majority of these pulsars did not favour timing models that included a $\ddot{\nu}$ component, and also show little evidence of a significant fractional evolution in \dot{P} as a function of P over our observing span.

[Figure 4.9](#) shows the values of P and \dot{P} for both the pulsars in our sample and those in [Parthasarathy et al. \(2019, 2020\)](#), where there is a rapid drop-off in the number of pulsars with significant $\ddot{\nu}$ measurements below a characteristic age of $\sim 10^5$ yr. Much of this can be put down to $\ddot{\nu}$ being harder to detect in pulsars with smaller values of ν and $\dot{\nu}$ (see [Equation 4.2](#)), where long timing baselines are needed to distinguish the resulting small fractional change in $\dot{\nu}$ from timing noise ([Parthasarathy et al., 2020](#)). However, there are some young, actively glitching pulsars among our sample that represent obvious outliers.

According to [Equation 4.8](#), all three of the most actively glitching pulsars in [Table 4.2](#) (PSRs J1341–6220, J1413–6141 and J1740–3015) should undergo a large glitch every 18, 11 and 60 yr respectively. These pulsars have not only undergone several large glitches in a much shorter span of time than their predicted glitch wait time, they have also exhibited a large number of small glitches (i.e. $\Delta\nu_g/\nu \lesssim 10^{-6}$) interspersed between them. Glitches of similar amplitudes are seldom seen among the Vela-like pulsars, and indeed Vela itself appears to have a real under-abundance of small amplitude glitches ([Howitt et al., 2018](#); [Espinoza et al., 2021](#)). A potential clue as to why these pulsars do not possess detectable

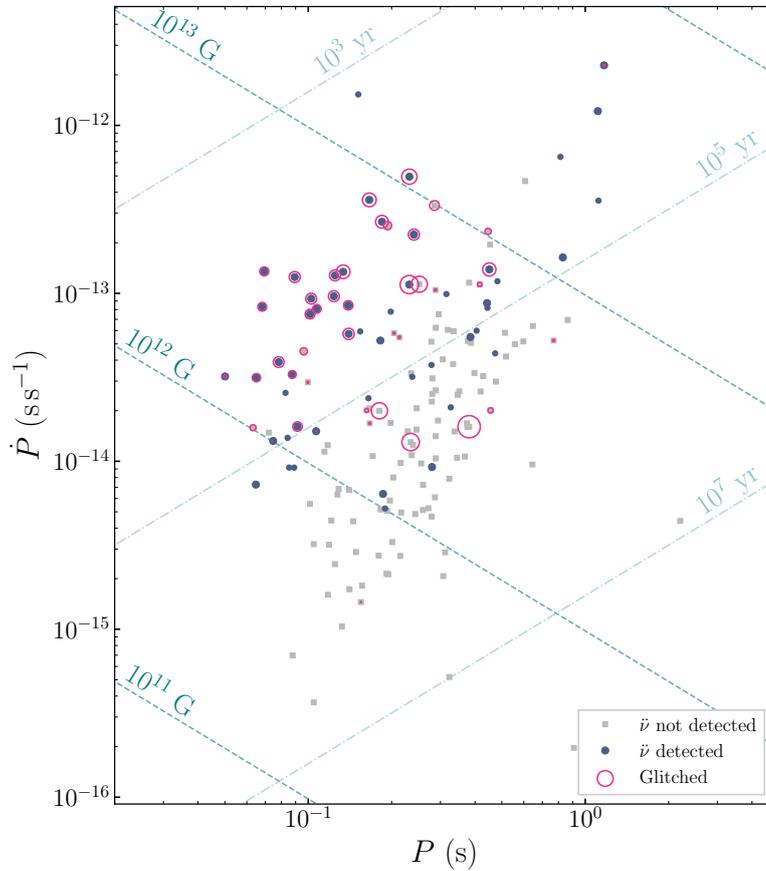


Figure 4.9 $P-\dot{P}$ diagram showing pulsars in both our sample and that of Parthasarathy et al. (2019, 2020) that favour (blue circles) or disfavour (grey squares) a timing model with a $\ddot{\nu}$ term. Pulsars that have glitched are highlighted with open circles, where the size is proportional to the amplitude of their largest glitch.

large n , despite being situated near the Vela-like population in $P-\dot{P}$ space, comes from the lack of consistently large $\Delta\dot{\nu}_g$ components associated with our glitch measurements in Table 4.2. Our strong observational relationship between n_{int} and $\Delta\dot{\nu}_g/T_g$ in Equation 4.13 suggests $\ddot{\nu}_{\text{int}} \rightarrow 0$ for $\Delta\dot{\nu}_g \rightarrow 0$ and $T_g \rightarrow \infty$. Hence, if there is genuinely no change in spin-down associated with these glitches, then there is no glitch recovery induced large- n process to detect in these pulsars. It is therefore possible that the underlying mechanism that drives the linear inter-glitch recoveries of the Vela-like pulsars is not active in these pulsars.

As noted in Section 4.5.2, the $P-\dot{P}$ post-glitch behaviour of the three pulsars with gigantic glitches (PSRs J1052–5954, J1650–4502 and J1757–2421) and PSR J0729–1448 is similar in that they all retained significant changes in their spin-down rates. A key

difference however, is that the pulsars with gigantic glitches all show evidence of at least one exponential recovery that causes some of their $\Delta\dot{\nu}_g$ to decay, whereas PSR J0729–1448 showed no evidence of a recovery. Such behaviour can arise if the core does not decouple from the crust during a glitch, but does so over long timescales after a glitch, leading to the appearance of a permanent increase in the spin-down rate (Haskell & Antonopoulou, 2014; Akbal et al., 2017). Our constraints on the braking index of PSR J0729–1448 returned $n = 2.3 \pm 0.9$ where the $\ddot{\nu}$ model is marginally disfavoured with $\ln(\mathcal{B}_{F2=0}^{F2 \neq 0}) = -1.3$, meaning its long-term braking index must be $n < 3.2$ (at 68% confidence) with no significant glitch-induced, large $\ddot{\nu}$ affecting its timing. As for why the pulsars with gigantic glitches retain a significant amount of their change in spin-down, it is possible they exhibit the same long-term linear recovery seen in the Vela-like pulsars, just on a much longer timescale than our current observations cover. Indeed the posterior distributions for n in these pulsars all peak away from zero, and for PSR J1052–5954 is well constrained to $n = 233 \pm 41$ albeit with a Bayes factor that marginally prefers a model containing a $\ddot{\nu}$ -term in addition to power-law red noise ($\ln(\mathcal{B}_{F2=0}^{F2 \neq 0}) = 2.7$), but does not exceed our $\ln(\mathcal{B}) > 3$ threshold. Hence longer-term observations of these pulsars may be able to distinguish their inter-glitch braking indices.

4.6 Conclusions

In this work we presented the results of a search for glitches in the timing of 74 young pulsars, along with inferences on the effects these glitches have on their rotational evolution. A total of 124 glitches were identified in these pulsars. The overall distributions of glitch properties from our sample largely reflects that of the broader population, and an analysis with the HMM-based glitch detection algorithm of Melatos et al. (2020) suggests our sample is complete to a pulsar averaged 90 percent upper-limit of $\Delta\nu_g^{90\%}/\nu \lesssim 8.1 \times 10^{-9}$. Having accounted for the glitches, we then conducted a Bayesian model selection study akin to that of Parthasarathy et al. (2019) for each of our pulsars. This led to the detection of a significant $\ddot{\nu}$ component, and subsequent measurement of the braking index, for 32 pulsars. Ten of these pulsars have never been seen to undergo a large amplitude glitch that contained a significant $\Delta\dot{\nu}_g$ -component, whereas the other pulsars exhibited glitches similar to those seen in the Vela pulsar. The measured values of n are uncorrelated with characteristic age, and the braking index distributions of non-glitching and glitching pulsars are indistinguishable from one another.

The observed $P-\dot{P}$ tracks of these Vela-like pulsars show that they evolve with a high inter-glitch braking index, n_{int} , and undergo a large change in \dot{P} at the time of the glitch.

The glitches seem to occur no later than a change in $\delta\dot{P}/\dot{P}$ of 0.004%. After accounting for the glitches, the decades-long evolution in $P\text{-}\dot{P}$ is consistent with a small n , similar to what was presented in [Espinoza et al. \(2017\)](#), in spite of the large n_{int} . We surmise that the sample of pulsars with large values of n in [Parthasarathy et al. \(2020\)](#) will undergo a large glitch in the near future and that they share similar characteristics to the pulsars in the present sample. We show there is a near one-to-one relationship between $\ddot{\nu}_{\text{int}}$ and $\Delta\dot{\nu}_g/T_g$ (Figure 4.8).

If the physical justification for this relationship is the vortex creep model proposed by [Alpar et al. \(1984b\)](#), then measurements of $\ddot{\nu}_{\text{int}}$ and $\Delta\dot{\nu}_g$ following a large glitch could be used to predict when the next large glitch may occur in these pulsars, as was done by [Akbal et al. \(2017\)](#) for the Vela pulsar. The use of an expanded version of this model that accounts for both over- and under-corrections in $\dot{\nu}$ between glitches (i.e. stochasticity in the fractional \dot{P} -jumps of PSRs J1301–6305, J1420–6048 and J1709–4429 in Figure 4.6) could provide useful forecasts for ongoing efforts to observe glitch events in real time with either dedicated search instruments near the glitch epoch (i.e. similar to that performed by [Dodson et al. 2002](#) and [Palfreyman et al. 2018](#) with the Mount Pleasant Observatory) or large scale pulsar monitoring programmes at UTMOST ([Jankowski et al., 2019](#); [Lower et al., 2020c](#)), CHIME/Pulsar ([CHIME/Pulsar Collaboration et al., 2021](#)) and the future SKA ([Watts et al., 2015](#); [Stappers et al., 2018](#)).

5

Spectropolarimetric properties of Swift J1818.0–1607: A 1.4 s radio magnetar

The soft-gamma repeater Swift J1818.0–1607 is only the fifth magnetar found to exhibit pulsed radio emission. Using the Ultra-Wideband Low receiver system of the Parkes radio telescope, we conducted a 3 h observation of Swift J1818.0–1607. Folding the data at a rotation period of $P = 1.363$ s, we obtained wideband polarization profiles and flux density measurements covering radio frequencies between 704 and 4032 MHz. After measuring, and then correcting for the pulsar’s rotation measure of $1442.0 \pm 0.2 \text{ rad m}^{-2}$, we find the radio profile is between 80-100 per cent linearly polarised across the wide observing band, with a small amount of depolarisation at low frequencies that we ascribe to scatter broadening. We also measure a steep spectral index of $\alpha = -2.26^{+0.02}_{-0.03}$ across our large frequency range, a significant deviation from the flat or inverted spectra often associated with radio-loud magnetars. The steep spectrum and temporal rise in flux density bears some resemblance to the behaviour of the magnetar-like, rotation-powered pulsar PSR J1119–6127. This leads us to speculate that Swift J1818.0–1607 may represent an additional link between rotation-powered pulsars and magnetars.

This chapter has been published in *The Astrophysical Journal Letters* (Lower et al., 2020a).

5.1 Introduction

Magnetars are a rare class of relatively slow rotating neutron star that are inferred to possess some of the strongest magnetic fields in the Universe. Until recently, only 4 of the 23 confirmed magnetars ¹ (Olausen & Kaspi, 2014) were seen to exhibit pulsed radio

¹<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>

emission (Camilo et al., 2006, 2007a; Levin et al., 2010; Eatough et al., 2013; Shannon & Johnston, 2013). Unlike standard rotation-powered pulsars, the radio pulses seen from these magnetars have generally flat spectra and display highly variable flux densities over timescales ranging between seconds to months (Camilo et al., 2007b; Lazaridis et al., 2008). Their single pulses are often comprised of many burst-like sub-pulses that display a remarkable range of temporal phenomenology. These sub-pulses have drawn comparisons to similar ‘spiky’ emission seen in high magnetic field strength pulsars (Weltevrede et al., 2011), rotating radio transients (RRATs; McLaughlin et al., 2006), and fast radio bursts (FRBs; Pearlman et al., 2018). Observations covering wide radio frequency bands may shed light on their magnetospheric conditions following outbursts, in particular whether the same processes that produce coherent, highly polarised emission in rotation-powered pulsars is also responsible for pulsed radio emission from magnetars.

Recently a fifth radio-bright magnetar was identified. *Swift* J1818.0–1607 was discovered by the *Swift* space observatory following the detection of a gamma-ray outburst by the Burst Alert Telescope on MJD 58920 (2020-03-12-21:16:47 UT). The burst was quickly localized to an X-ray point source by the on-board X-ray telescope (Evans et al., 2020). Observations by the *Neutron star Interior Composition Explorer* found the source exhibited pulsed X-ray emission with a periodicity of 1.36 s (Enoto et al., 2020). Two days after the initial outburst, highly linearly polarised radio pulsations were detected with a dispersion measure (DM) of $706 \pm 4 \text{ pc cm}^{-3}$ during follow-up observations by the 100-m Effelsberg radio telescope observing in a band centered on 1.37 GHz (Karuppusamy et al., 2020). Continued timing provided an initial measurement of the spin-period derivative, $\dot{P} = 9 \pm 1 \times 10^{-11}$ (Esposito et al., 2020), firmly cementing its status as the fastest rotating, and possibly the youngest magnetar found to date. Observations performed at multiple radio wavelengths indicated the magnetar’s radio emission has a steep spectral index (Gajjar et al., 2020; Lower & Shannon, 2020). This is similar to the observed radio spectra of many ordinary, rotation-powered radio pulsars, but significantly differs from the flat or inverted spectra of the four other radio loud magnetars. The apparently low surface temperature (Esposito et al., 2020) and lack of coincident supernova remnant, indicate *Swift* J1818.0–1607 may be significantly older than implied by its characteristic age of 240 – 310 yrs, and may represent a transitional link between magnetars and the population of high B-field, rotation-powered pulsars. In particular, the reported spectral flattening by Majid et al. (2020b) may indicate a possible link to the 2016 magnetar-like outburst of PSR J1119–6127 (Majid et al., 2017).

In this letter we report on observations of *Swift* J1818.0–1607 using the Ultra-wideband

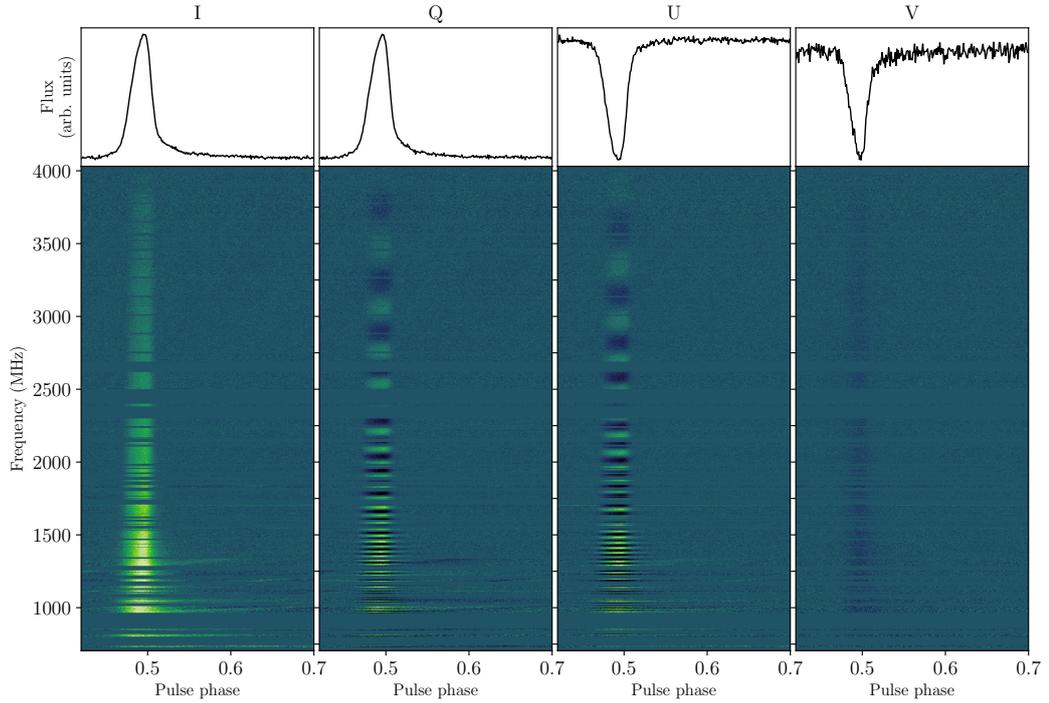


Figure 5.1 Faraday-corrected average polarization profiles (*top*) and uncorrected, time-averaged polarization spectra (*bottom*) of Swift J1818.0–1607. All four Stokes parameters are plotted with 2 MHz spectral resolution and 0.67 ms temporal resolution. The large rotation measure of $1442.0 \pm 0.2 \text{ rad m}^{-2}$ is clearly visible in Stokes Q and U. Horizontal gaps in each panel represent frequency channels that were excised due to RFI contamination. Some broadband sweeps of RFI remain visible below 1300 MHz.

Low (UWL) receiver system (Hobbs et al., 2020) of the CSIRO 64-m Parkes radio telescope. Using Bayesian inference techniques, we measured the broadband properties of the time averaged polarization spectrum and analysed the sample of bright single pulses observed throughout the approximately 3 hour-long observation. We then compare these results to previous observations of the four other radio loud magnetars and the general pulsar population. Finally, we discuss the potential evolutionary pathways of Swift J1818.0–1607.

5.2 Observation and analysis

We conducted a 10473 s observation of Swift J1818.0–1607 on MJD 58939 using the Parkes UWL receiver (Hobbs et al., 2020) under the target of opportunity request PX057 (PI: Lower). Pulsar search-mode data with $128 \mu\text{s}$ sampling covering the full UWL band from 704–4032 MHz with full Stokes information were recorded using the MEDUSA backend and coherently dedispersed on a channel by channel basis at a DM of 700 pc cm^{-3} to minimize

Table 5.1 Scatter broadening (τ_{sc}), period-averaged flux density (S_ν) measurements, and fractional linear ($\langle L/I \rangle$) and circular ($\langle |V|/I \rangle$) polarization of each 256 MHz sub-band.

Frequency (MHz)	τ_{sc} (ms)	S_ν (mJy)	$\langle L/I \rangle$	$\langle V /I \rangle$
3879	$\lesssim 3$	0.31 ± 0.03	0.73	0.19
3656	$\lesssim 3$	0.33 ± 0.01	0.94	0.18
3386	$\lesssim 3$	0.41 ± 0.01	0.86	0.18
3137	$\lesssim 1.9$	0.52 ± 0.01	0.88	0.20
2880	0.8 ± 0.5	0.62 ± 0.01	0.88	0.18
2612	1.2 ± 0.6	0.82 ± 0.01	0.92	0.16
2304	2.8 ± 0.5	1.11 ± 0.02	0.73	0.12
2106	3.5 ± 0.2	1.40 ± 0.01	0.92	0.11
1858	5.3 ± 0.2	1.79 ± 0.02	0.97	0.12
1598	8.8 ± 0.2	2.53 ± 0.01	0.93	0.11
1356	16.8 ± 0.2	3.72 ± 0.1	0.93	0.10
1070	38.6 ± 0.5	6.0 ± 0.1	0.83	0.16
809	186^{+7}_{-6}	11.8 ± 0.6	0.52	0.18

Note: The uncertainties denote the 68% confidence intervals. Only upper limits are set on the scattering timescale at frequencies above 2880 MHz and are with 68% confidence.

dispersive smearing of the pulse profile. Note the profiles shown in Figure 5.1 have been dedispersed using the inferred DM of 706.0 pc cm^{-3} from Section 5.2.3. The data were then folded at the pulse period of the magnetar using DSPSR (van Straten & Bailes, 2011) and saved to a `psrfits` (Hotan et al., 2004) format archive with 1024 phase bins, and 3328 frequency channels with 1 MHz frequency resolution. Approximately 35 per cent of the 3328 frequency channels were heavily contaminated by radio frequency interference (RFI), and were subsequently excised using the standard `paz` and `pazi` tools in PSRCHIVE (Hotan et al., 2004; van Straten et al., 2012). The data were flux and polarization calibrated in the same manner as Dai et al. (2019), with the exception that we used the radio galaxy PKS B0407–658 as a flux density reference instead of 3C 218. Unlike 3C 218, PKS B0407–658 is not resolved by Parkes above $\sim 3 \text{ GHz}$, making it a more reliable calibrator for the UWL. We used an observation of a linearly polarised noise diode prior to observing the magnetar, in addition to on- and off-source observations of PKS B0407–658 taken on MJD 58638 to measure the noise diode brightness and to correct the phase and absolute gain of the system. We note that any leakage terms were not corrected for, which may be of order 5 per cent toward the top of the band.

5.2.1 Profile phenomenology and flux density

Dynamic spectra showing the four Stokes parameters across the continuous 704-4032 MHz UWL band are displayed in Figure 5.1. The pulse profile shows clear evidence of a steep negative gradient in flux density, and can be described as the superposition of two Gaussian components (G_1 and G_2 hereafter). The narrower G_2 component appears brighter toward the lower end of the UWL band, indicating it has a steeper spectral index than G_1 .

We further analysed the profile by dividing the data into 13 sub-bands, each having 256 MHz of bandwidth. These sub-bands were then averaged in frequency and polarization before being fit with a two-component Gaussian profile convolved with a one-sided exponential pulse broadening function

$$f(t) = \sum_{i=1}^2 \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-(t-\mu_i)^2/2\sigma_i^2} \otimes e^{-t/\tau_{\text{sc}}}, \quad (5.1)$$

where μ_i and σ_i are the centroids and widths of the i -th Gaussian component, \otimes indicates a convolution and τ_{sc} is the scattering timescale. The resulting posterior probability distributions were sampled using the `bilby` software package (Ashton et al., 2019a) as a front-end to the `dynesty` nested sampling algorithm (Speagle, 2020). Initially we fit the sub-bands assuming uniform priors on the widths of the profile components G_1 and G_2 . However, we found the component widths were highly covariant with the scattering timescale, to the point where we could only recover upper-limits for scattering in sub-bands above 2106 MHz. As the profile width does not appear to undergo significant evolution with frequency, aside from scatter broadening, we re-fit the sub-banded data assuming Gaussian priors of $\pi(\sigma_1) \sim \mathcal{N}(8 \text{ ms}, 1 \text{ ms})$ and $\pi(\sigma_2) \sim \mathcal{N}(7 \text{ ms}, 1 \text{ ms})$ for the widths of G_1 and G_2 respectively.

The resulting scattering timescale and period-averaged flux density – measured by averaging the best-fit template for each sub-band in pulse phase – are presented in Table 5.1. We measure a scattering timescale referenced to 1 GHz of $\tau_{\text{sc},1 \text{ GHz}} = 42_{-3}^{+9}$ ms, with a scattering index of $\alpha_{\text{sc}} = -3.4_{-0.2}^{+0.3}$. Similar but less well constrained values of $\alpha_{\text{sc}} = -3.6_{-1.1}^{+0.8}$ and $\tau_{\text{sc},1 \text{ GHz}} = 41_{-18}^{+19}$ ms were obtained when we used uniform priors on the widths of G_1 and G_2 . In either case, the scattering timescale is consistent with the expected value of 62 ± 30 ms from the NE2001 galactic electron density model at 1 GHz (Cordes & Lazio, 2002). While the scattering index is smaller than the expected value of $\alpha_{\text{sc}} = -4$ or $\alpha_{\text{sc}} = -4.4$ expected from Kolmogorov turbulence, they are consistent with the scattering indices of many other pulsars (see for example Geyer et al., 2017). We also fit the period-

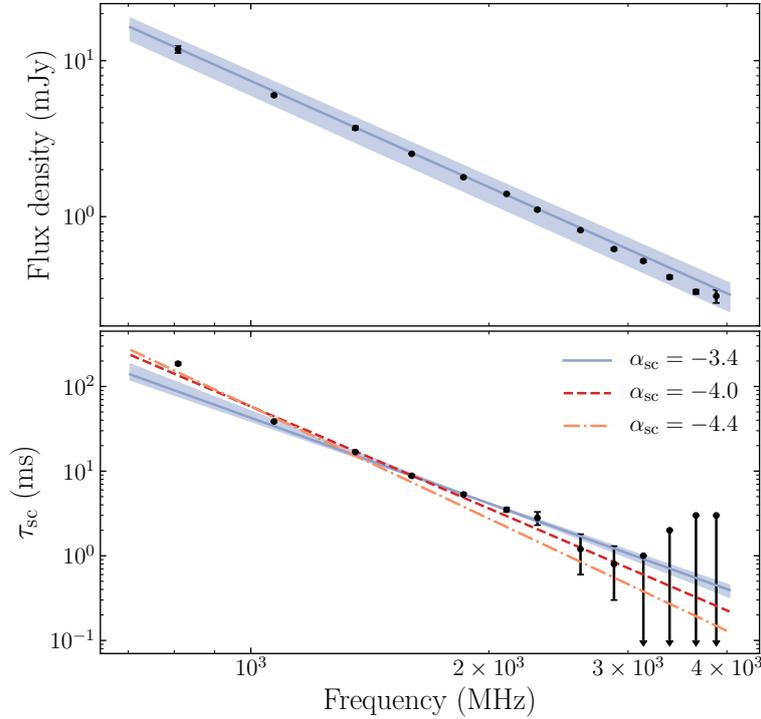


Figure 5.2 Period-averaged flux density (*top*) and scattering timescale (*bottom*) as a functions of frequency. The blue solid lines indicates the median fit while the shaded region is bounded by the 68% confidence intervals. Dashed red and dash-dotted orange lines correspond to scattering indices of -4 and -4.4 respectively.

averaged flux density spectrum using a simple power-law function, $S_\nu \propto \nu^\alpha$, obtaining a spectral index of $\alpha = -2.26^{+0.02}_{-0.03}$. The fits to the spectral index and scattering timescale are plotted in Figure 5.2. The reduced χ^2 for the scattering relation shown in Figure 5.2 is 13.8. We attribute the high value to overestimation of the scattering timescale in the RFI-affected 809 MHz band. Removing the 809 MHz data point confirms this suspicion, as refitting the scattering relation returns a consistent scattering index of $\alpha = -3.6^{+0.4}_{-0.3}$ and a reduced χ^2 of 0.6.

5.2.2 Polarimetry

Figure 5.1 clearly shows the linear polarization has undergone significant Faraday rotation, as evidenced by the large number of changes in sign for Stokes Q and U . Following the Bayesian methodology presented in Bannister et al. (2019), we measured the phase averaged rotation measure (RM) of the magnetar by directly fitting Stokes Q and U as

a function of frequency, obtaining a value of $1442.0 \pm 0.2 \text{ rad m}^{-2}$ (68 per cent confidence interval). Note, this measurement does not include corrections for the ionosphere which can often exceed our measurement uncertainty. At Parkes, the ionospheric contribution is typically between -0.2 to -2.0 rad m^{-2} (Han et al., 2018).

To better visualize the polarization profiles, we plot the averaged polarization pulse profiles at 13 frequencies in Figure 5.3, along with the linear polarization position angle for each sub-band. The pulse profile is more than 90 per cent linearly polarised across most of the UWL band, although a small amount of circular polarization is also present. Apparent depolarisation due to scatter broadening (Li & Han, 2003) is evident below 1356 MHz. Slight variations in the fractional linear and circular polarizations listed in Table 5.1 likely result from a combination of noise and polarization impurities in the receiver system. The apparent depolarisation in the 2304 MHz band is an artefact of residual RFI from wireless communications contaminating the narrow strip of non-excised channels between 2380 and 2400 MHz. Additionally, the lack of polarisation in the ‘bump’ visible in the off-pulse noise of the 3879 MHz sub-band suggests this feature is likely to be residual impulsive RFI, not an additional profile component. There is a slight upward slope in the linear polarization position angle (PA), with little frequency dependent evolution except for scatter-induced smearing at lower frequencies.

5.2.3 Single pulses

To analyse the single pulses from the magnetar, we created single pulse archives from the original `psrfits` search-mode filterbank. We then performed a boxcar search for single pulses on copies of these archives where all frequency channels outside the 1300-2500 MHz band had been excised to minimize confusion with RFI. We limited this search to only the on-pulse region of each archive. Applying a maximum boxcar width of 85 ms and threshold S/N of 7, we find 5052 of the 7008 single pulse archives contained a single pulse candidate that met our criterion with a median S/N of 13.8. Upon visual inspection, we found the single pulses typically consist of 1-3 ‘spiky’ sub-pulses with similar phenomenology to single pulses seen from the four other radio loud magnetars. We did not observe any single pulses emitted at rotational phases outside the ‘on-pulse’ region represented by the integrated profiles in Figure 5.3, nor evidence of sporadic pulses from the additional profile component reported by Maan & van Leeuwen (2020). Occasional gaps or nulls in emission were seen throughout the observation. Similar behaviour has been reported in observations of the galactic centre magnetar SGR J1745–2900 (Yan et al., 2018). However, it is not clear whether the gaps we observed represent true nulls, where the radio emission

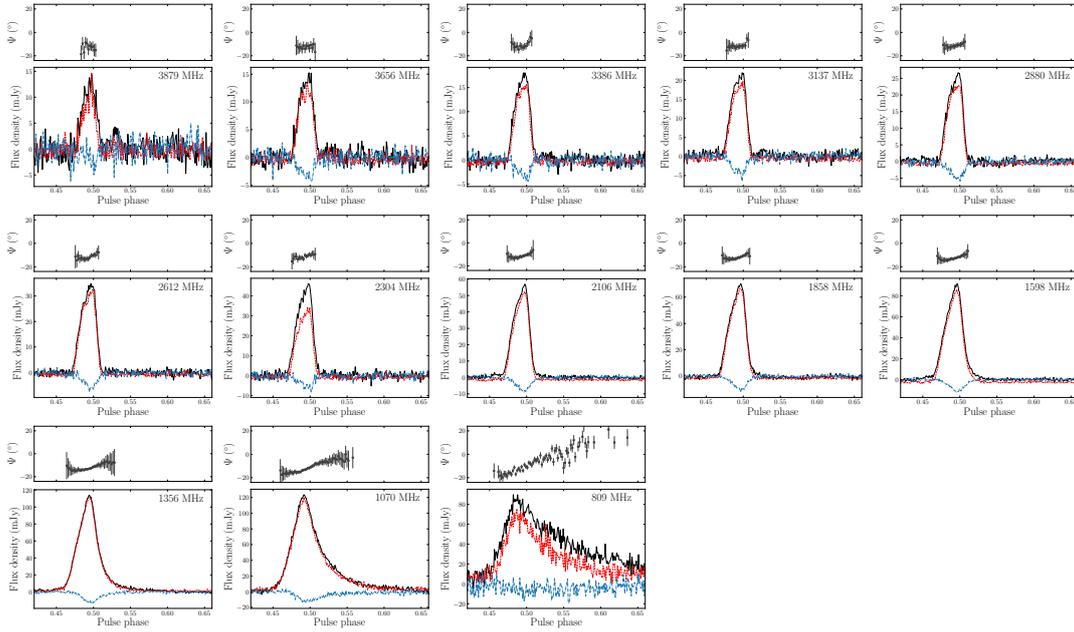


Figure 5.3 Polarization profiles of Swift J1818.0–1607 averaged over 13 sub-bands from 3656 MHz to 809 MHz, each dedispersed at a DM of 706 pc cm^{-3} and covering 256 MHz of bandwidth. Black represents total intensity, red linear polarization and blue is circular polarization. The linear polarization position angles (Ψ) are corrected for the $\text{RM} = 1442.0 \pm 0.2 \text{ rad m}^{-2}$ at a reference frequency of 2368 MHz.

mechanism completely shuts off, or if the radio pulses during these rotations were simply below the detection threshold of the receiver.

We measured the flux density of the on- and off-pulse regions of each single-pulse archive using the `psrflux` tool from PSRCHIVE by cross-correlating the data with a scatter-broadened Gaussian template. Both the on- and off-pulse flux density measurements were then converted to units of matched-filter S/N by scaling each measurement by a factor of 1.4 – the scale factor needed to scale the off-pulse distribution such that it has a mean of zero and variance of one. The resulting on- and off-pulse S/N distributions are shown in Figure 5.4. We note this definition of S/N is different to the one used in the earlier single pulse search, which was a top-hat S/N used to place quantitative constraints on the number of single pulses we detected. Negative S/N ratios can be attributed to the on-pulse flux being below zero due to fluctuations in the baseline. The on-pulse distribution is well described by a log-normal with a log-mean of 1.925 ± 0.003 and width of 0.25 ± 0.01 that has been convolved with a Gaussian distribution with zero mean and unit variance. This distribution width is typical of the rotation-powered pulsar population as a whole (Burke-Spolaor et al., 2012). While there are some outliers, the lack of a power-

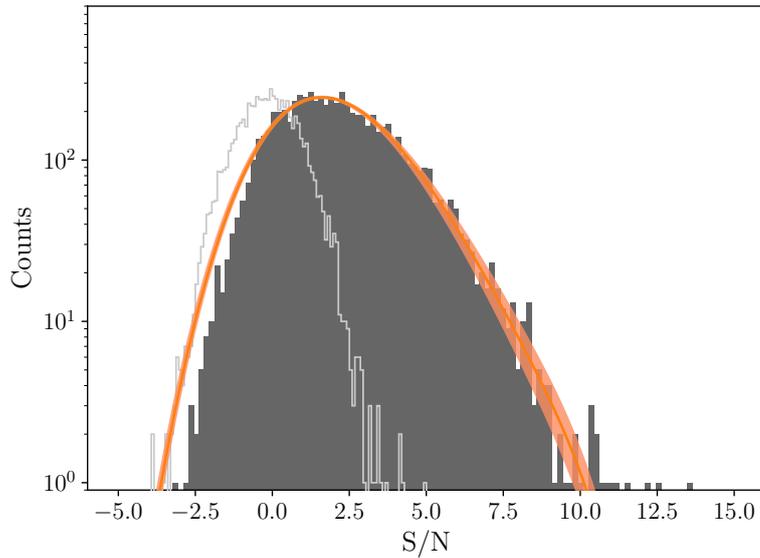


Figure 5.4 Matched-filter S/N distribution for the frequency-averaged single pulses (dark grey), scaled such that the off-pulse noise (light grey) has zero mean and unit variance. The orange line and shading are the median log-normal convolved with a Gaussian fit to the data and associated 68% confidence intervals.

law tail in the distribution indicates no giant pulses were detected during our observation, contradictory to the claim by [Esposito et al. \(2020\)](#) that the single pulses are dominated by sporadic giant pulses. It is possible their giant pulse detections originated from the transient profile component seen in early observations by [Maan & van Leeuwen \(2020\)](#), which had disappeared sometime prior to our observation with Parkes.

The narrow widths of magnetar single pulses and sub-pulses enable high-accuracy DM measurements, particularly when observed across large bandwidths. For example, the bright single pulse shown in [Figure 5.5](#) returned a structure-optimized DM of $707.3 \pm 0.2 \text{ pc cm}^{-3}$. Repeating this for the brightest 215 single pulses in our sample, we find the distribution of structure-optimised DMs is well described by a Gaussian with a mean of 706.0 pc cm^{-3} and a standard deviation of 2.6 pc cm^{-3} . From this, we estimated the magnetar’s DM to be $706.0 \pm 0.2 \text{ pc cm}^{-3}$ where the uncertainty is derived from the standard deviation of the DM distribution $\sigma_{\text{DM}} = 2.6/(215 - 2)^{1/2} \text{ pc cm}^{-3}$. The variations in DM are more likely to have resulted from systematic errors in the structure-optimization algorithm combined with the variable number of sub-pulses in each pulse as opposed to short-timescale variations in the local environment of the magnetar. Long-term monitoring over year-long timescales will reveal if Swift J1818.0–1607 experiences DM variations similar to those seen in repeating FRBs (e.g. [Hessels et al., 2019](#)).

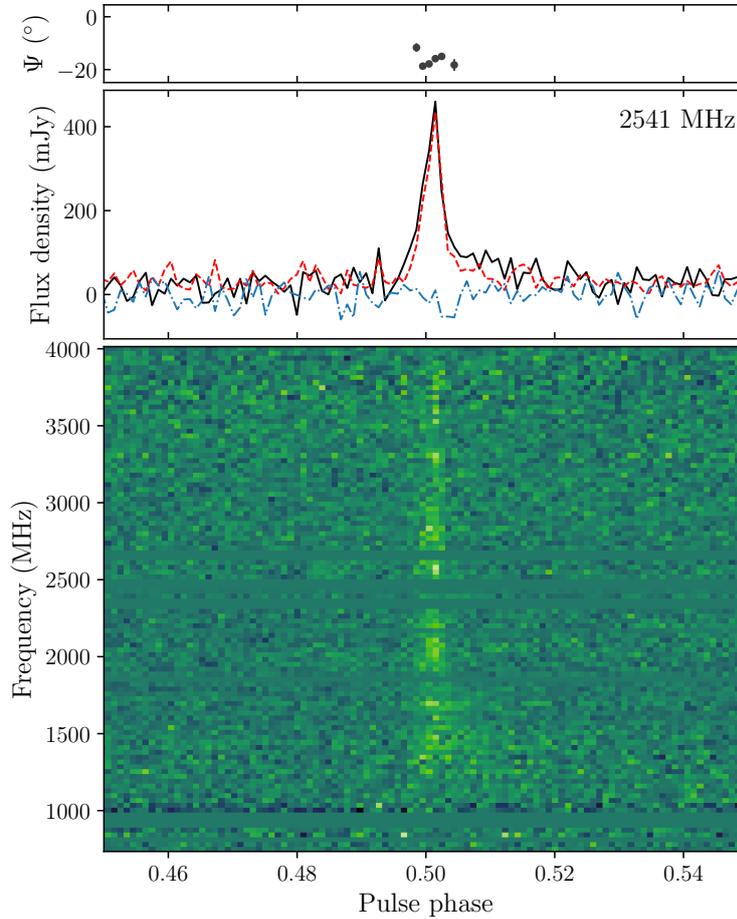


Figure 5.5 A single pulse from Swift J1818.0–1607. The top and middle panels show the position angle and integrated polarization profile. The bottom panel shows the waterfall diagram of the pulse dedispersed at a $DM = 707.3 \pm 0.2 \text{ pc cm}^{-3}$ with 0.67 ms time resolution and 16 MHz spectral resolution.

Using the NE2001 (Cordes & Lazio, 2002) and YMW16 (Yao et al., 2017) galactic free electron density models, the distance to the magnetar is estimated to be either $8.1 \pm 1.6 \text{ kpc}$ (NE2001) or 4.8 kpc (YMW16), where the uncertainty is dominated by the model chosen.

From our measurements of both the RM and DM, we can estimate the average parallel magnetic field strength along the line of sight to the magnetar using the equation $B_{\parallel} = 1.2\text{RM}/\text{DM}$, where B_{\parallel} is in units of μG , and the RM and DM are in their usual units (rad m^{-2} and pc cm^{-3}). Our measured value of $2.5 \mu\text{G}$ is fairly typical of line-of-sight B_{\parallel} measurements from pulsars within the galactic plane (Han et al., 2018)

5.3 Discussion

In general, the pulsed radio emission from Swift J1818.0–1607 shares a lot of the same phenomenology seen in other radio loud magnetars: a high degree of linear polarization, burst-like sub-pulses and extremely variable pulse-to-pulse flux densities. However, the steep spectral index we measure is more consistent with the spectral indices of many rotation powered pulsars when compared to the flat spectral indices of the four other radio magnetars which typically range between -0.5 to $+0.3$ (Lazaridis et al., 2008; Torne et al., 2015; Dai et al., 2019), making this new magnetar a significant outlier. Given the DM and location of the magnetar, the effects of diffractive interstellar scintillation are negligible at the UWL observing band. For instance, the NE2001 model predicts a scintillation bandwidth of only 3_{-1}^{+3} Hz at 1 GHz. Hence the steep spectrum is intrinsic to Swift J1818.0–1607. This indicates that it was premature to assume that all radio magnetars have flat spectra. At the large DMs typical of magnetars, those that have steep radio spectra might be so scatter-broadened as to induce a significant selection effect towards those with flatter spectra. When compared to the 276 pulsars in Jankowski et al. (2018) that have spectra best fit by a simple power-law, only ~ 11 per cent of pulsars have steeper spectra than Swift J1818.0–1607, while the four other radio magnetars all have spectral indices that are flatter than ~ 94 per cent of their sample. Hence, Swift J1818.0–1607 may be an example of the diversity that could exist in the wider, as-of-yet undetected radio magnetar population. The spectral properties could also be related to the magnetar possessing a less evolved magnetic field structure due to its youth.

Assuming Swift J1818.0–1607 was born rapidly rotating ($P \sim 10$ ms) and its spin-down is dominated by magnetic dipole radiation (braking index = 3), measurements of its spin and spin-down place its characteristic age between only 240–310 yrs (Champion et al., 2020b; Hu et al., 2020a; Esposito et al., 2020), the second smallest of any pulsar after SGR J1806–20 (Mereghetti et al., 2005). However, given large amount of uncertainty surrounding neutron star rotation periods at birth and the diversity in measured pulsar (and magnetar) braking indices, its true age is likely to be significantly different than the inferred spin-down age. Indeed the period derivatives of magnetars can change by large factors within just a few years (see for example Scholz et al., 2017). A more accurate kinematic age could be inferred from associating the magnetar to a progenitor supernova remnant, combined with a proper-motion measurement from very-long baseline interferometry. However, we find there are no catalogued supernova remnants or pulsar-wind nebula co-located with its position (Green, 2019). The two closest supernova remnants (G014.3+0.1 and G014.1–0.1) are approximately 19 arcmin and 27 arcmin away from the

position of the magnetar on sky (Galactic coordinates: $l = 14.8^\circ$, $b = -0.14^\circ$) respectively, making an association highly unlikely. The lack of an associated supernova remnant is not too surprising, as only eight of the twenty-three known magnetars have claimed associations. Additionally, the strong spin-down powered wind from new-born magnetars can accelerate the remnant expansion to the point that only anomalously diffuse shells, or no remnant at all, remains on century-long timescales (Duncan & Thompson, 1992). If the progenitor supernova remnant has not been dissipated, then deep radio and X-ray imaging may be able to detect it.

Alternatively, we speculate the steep spectrum and its unusually faint X-ray luminosity of 7×10^{34} ergs s⁻¹ (Esposito et al., 2020)² may be evidence this new magnetar was initially born as a rotation powered pulsar that obtained the rotational properties of a magnetar over time, similar to what is predicted for PSR J1734–3333 (Espinoza et al., 2011a). Such evolution can occur if the magnetic and spin axes underwent rapid alignment over time (Johnston & Karastergiou, 2017), or if the pulsar underwent an extended period of magnetic field growth after the surface magnetic field was initially buried due to fall-back accretion (e.g. Ho et al., 2015).

If the properties of Swift J1818.0–1607 are the result of rapid magnetic and spin axes alignment, we would expect the PA to be consistent with that of an aligned rotator. There is some evidence magnetars tend toward aligned spin and magnetic axes. Both 1E 1547.0–5408 and PSR J1622–4950 have PA swings that are consistent with being aligned rotators (Camilo et al., 2008; Levin et al., 2012). This is further backed up by the wide radio profiles, and low pulsed X-ray fractions of these two magnetars (Halpern et al., 2008; Camilo et al., 2018). There is some ambiguity as to whether the spin and magnetic axes of XTE J1810–197 are aligned or orthogonal, as Camilo et al. (2007b) found both scenarios adequately describe the PA swing across its main pulse and inter-pulse. Conversely, Kramer et al. (2007) found that an offset dipole described by two separate rotating vector models (RVMs, Radhakrishnan & Cooke, 1969) could also describe its PA behaviour, and speculated it may be evidence for XTE J1810–197 having a multipole magnetic field. Additionally, Dai et al. (2019) observed distinctly non-RVM PA variations following its 2018 outburst. For Swift J1818.0–1607, the flat PA in the higher-frequency panels of Figure 5.3 is broadly consistent with the RVM for a dipole magnetic field. However, the narrow pulse duty cycle makes it difficult to constrain the star’s magnetic geometry, as the relatively flat PA could be consistent with either nearly aligned

²As noted in Esposito et al. (2020), the quoted X-ray luminosity assumes the smaller, YMW16 DM distance to the magnetar, and that a larger source distance (as implied by the NE2001 model) may yield a more normal luminosity.

magnetic and spin axes, or a large offset between the magnetic axis and our line-of-sight. Given the radio profiles of magnetars evolve over the weeks to months following an outburst (Kramer et al., 2007; Dai et al., 2019), it may be possible to measure the magnetic geometry of Swift J1818.0–1607 in the future.

Pulsars that experienced fall-back accretion soon after their birth can undergo apparent magnetic field growth as their magnetic fields diffuse to the surface over time (see for example Muslimov & Page, 1995). This can result in a seemingly ‘normal’ rotation-powered, young pulsar obtain magnetar-like rotational properties within $\sim 1\text{--}10$ kyr (Ho et al., 2015). If Swift J1818.0–1607 is a result of this evolutionary path, then we may expect it to show similar radio properties to the high B-field PSRs J1119–6127, J1208–6238 and J1846–0258. While PSRs J1846–0258 (Gavriil et al., 2008) and J1119–6127 (Archibald et al., 2016) have been observed to undergo magnetar-like outbursts in the past, only PSR J1119–6127 has been observed to emit radio pulses. Observationally, we can draw parallels between the radio properties of Swift J1818.0–1607 and those of PSR J1119–6127 during its 2016 outburst. Following the initial suppression and re-emergence of radio pulses from PSR J1119–6127, multi-band flux measurements found the pulsar possessed a steeper radio spectrum than its nominal $\alpha = -1.4 \pm 0.2$, with values of α ranging between -2.2 ± 0.2 to -1.9 ± 0.2 (Majid et al., 2017). Later observations found its radio spectrum had undergone spectral flattening to a more magnetar-like spectral index of -0.52 ± 0.06 over the months following the outburst (Pearlman et al., 2016). The flux density of PSR J1119–6127 also underwent a factor of 5 increase in two weeks after the outburst before recovering back to its normal levels (Dai et al., 2018). In addition to having a comparably steep post-outburst spectral index, Swift J1818.0–1607 appears to have also undergone a similar radio brightening, as the flux densities at 1356 MHz and 1598 MHz in Table 5.1 are a factor of 5–12 times higher than measurements at similar observing frequencies two weeks prior to our Parkes UWL observation (Karuppusamy et al., 2020; Esposito et al., 2020; Lower & Shannon, 2020). The refractive modulation timescale is expected to be very long (years) and the modulation index to be low (Cordes & Lazio, 2002). Thus the increase in flux density cannot be ascribed to refractive effects. If the current outburst of Swift J1818.0–1607 continues to proceed in a similar manner to the 2016 outburst of PSR J1119–6127, then we may expect the steep spectral index to undergo a similar flattening and for the flux density to decay to a more steady state over the coming months. A more recent spectral index measurement of $\alpha = -1.9 \pm 0.2$ from multi-band observations (Majid et al., 2020b) suggests some amount of spectral-flattening may have already occurred. Continued monitoring with multi-band and wide-bandwidth receiver systems will either

confirm the spectral index is flattening toward a more magnetar-like value, or is simply fluctuating about some mean value. Additionally, a measurement of the braking index would allow us to understand the future spin and magnetic field evolution of the magnetar and potentially confirm or rule out a rotation-powered pulsar origin.

6

The dynamic magnetosphere of Swift J1818.0–1607

Radio-loud magnetars display a wide variety of radio-pulse phenomenology seldom seen among the population of rotation-powered pulsars. Spectropolarimetry of the radio pulses from these objects has the potential to place constraints on their magnetic topology and unveil clues about the magnetar radio emission mechanism. Here we report on eight observations of the magnetar Swift J1818.0–1607 taken with the Parkes Ultra-Wideband Low receiver covering a wide frequency range from 0.7 to 4 GHz over a period of 5 months. The magnetar exhibits significant temporal profile evolution over this period, including the emergence of a new profile component with an inverted spectrum, two distinct types of radio emission mode switching, detected during two separate observations, and the appearance and disappearance of multiple polarization modes. These various phenomena are likely a result of ongoing reconfiguration of the plasma content and electric currents within the magnetosphere. Geometric fits to the linearly polarized position angle indicate we are viewing the magnetar at an angle of $\sim 99^\circ$ from the spin axis, and its magnetic and rotation axes are misaligned by $\sim 112^\circ$. While conducting these fits, we found the position angle swing had reversed direction on MJD 59062 compared to observations taken 15 days earlier and 12 days later. We speculate this phenomena may be evidence the radio emission from this magnetar originates from magnetic field lines associated with two co-located magnetic poles that are connected by a coronal loop.

This chapter was published in full in *Monthly Notices of the Royal Astronomical Society* (Lower et al., 2021)

6.1 Introduction

Swift J1818.0–1607 belongs to a sub-class of slowly rotating, young neutron stars that possess unusually high X-ray and gamma-ray luminosities, commonly referred to as magnetars. They are believed to be powered by the dissipation of their ultra-strong internal magnetic fields as opposed to the slow release of stored angular momentum (Thompson & Duncan, 1995). Most are detected as persistent sources of high energy electromagnetic radiation and occasionally undergo periods of high activity, where bursts of intense X-ray and gamma-ray emission are commonplace. If a magnetar was born rapidly rotating, its internal magnetic field will be strongly wound up (Duncan & Thompson, 1992). Relaxation of the internal magnetic field exerts strong magnetic forces on the crust that can lead to local or even global twists in the magnetic field due to horizontal plastic deformation or fracturing of the crust (i.e a starquake) if these stresses are allowed to build up over time (Thompson et al., 2002). It is the sudden twisting of the magnetic field lines along with magnetic re-connection events that are believed to power magnetar outbursts (see Kaspi & Beloborodov, 2017, for a review).

Beloborodov (2009) showed the current bundles that flow along a twist near the dipole axis of the magnetosphere can generate the conditions required for coherent radio emission to take place, potentially explaining why a handful of active magnetars have now been detected as radio pulsars. These ‘radio-loud’ magnetars exhibit an extremely diverse variety of radio emission phenomenology that are rarely displayed by less magnetic rotation-powered pulsars. Both their average and single pulse profiles have high degrees of linear polarization, typically in excess of 90 per cent (Kramer et al., 2007; Camilo et al., 2007a; Levin et al., 2010; Eatough et al., 2013), and often possess extremely flat radio spectra (Levin et al., 2012; Torne et al., 2015; Dai et al., 2019). The untwisting of their dynamic magnetic fields and associated electric currents following an outburst are imprinted in their radio profiles, which show variations in intensity and polarization, along with the emergence or disappearance of profile components on timescales ranging from a few hours to many months (e.g. Camilo et al., 2007c, 2016; Scholz et al., 2017). The sweep of the linear polarization position angle can be interpreted geometrically, as has been done for several of the radio magnetars (Camilo et al., 2007d,a; Levin et al., 2012). However, deviations from the standard models often employed to fit the position angle swing have led some to speculate on the role of emission from closed magnetic field lines and contributions from higher-order multipole magnetic fields (e.g. Kramer et al., 2007). Despite these deviations from standard pulsar behaviour, radio-loud magnetars generally have higher spin-down luminosities than most ‘radio-quiet’ magnetars, potentially point-

ing to a strong relationship with young radio pulsars (see [Rea et al. 2012](#) and discussions therein).

Unlike standard radio pulsars, the single pulses detected from magnetars are typically comprised of many ‘spiky’ sub-pulses that show highly variability in intensity and width on a pulse-to-pulse basis ([Serylak et al., 2009](#); [Levin et al., 2012](#); [Pearlman et al., 2018](#)). Similarities between the single pulse properties of magnetars and the phenomenology of fast radio bursts (FRBs; e.g. [Pearlman et al., 2018](#); [Maan et al., 2019](#)), combined with numerous FRB progenitor theories that invoke a magnetar central engine tentatively indicate radio magnetars within the Milky-Way may be galactic analogues to FRB progenitors (e.g. [Wadiasingh & Timokhin, 2019](#)). This possible connection has been strengthened by the detection of an extremely luminous, millisecond-duration radio burst from SGR 1935+2154 by the CHIME/FRB and STARE2 experiments ([CHIME/FRB Collaboration et al., 2020](#); [Bochenek et al., 2020](#)).

With a spin-period of approximately 1.4 s, Swift J1818.0–1607 is among the fastest rotating pulsars that show magnetar-like activity. A secular spin-down rate of $4.6 \times 10^{-11} \text{ s s}^{-1}$ and an inferred surface dipole magnetic field strength of $2.5 \times 10^{14} \text{ G}$ ([Champion et al., 2020a](#)), place Swift J1818.0–1607 among the growing population of known galactic magnetars ([Olausen & Kaspi, 2014](#))¹. Observations by the Effelsberg and Lovell radio telescopes soon after its discovery revealed the magnetar to be radio-bright ([Champion et al., 2020a](#)), making it only the fifth radio-loud magnetar.

Like other radio magnetars, its single pulses are comprised of narrow, spiky sub-pulses (see Figure 3 of [Esposito et al., 2020](#)), with a high degree of linear polarization across a wide range of frequencies ([Lower et al., 2020c](#)). However, its unusually steep radio spectrum and lower than anticipated quiescent X-ray luminosity ([Esposito et al., 2020](#)) seem to imply it shares more in common with more ordinary rotation-powered pulsars than other radio-loud magnetars. These irregular properties and similar behaviour to that of PSR J1119–6127 following its 2016 outburst ([Archibald et al., 2016](#); [Dai et al., 2018](#)) led to speculation that Swift J1818.0–1607 may represent a possible missing link between magnetars and the population of magnetar-like, high magnetic field strength (high B-field) pulsars ([Hu et al., 2020b](#)).

In this work, we explore the spectral, temporal and polarimetric properties of Swift J1818.0–1607 across the 3.3 GHz bandwidth of the Ultra-Wideband Low (UWL) receiver system of the CSIRO Parkes 64-m radio telescope (also known as *Murriyang*), covering eight epochs after its discovery in March 2020 until October 2020. The details of our

¹<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>

observations along with the calibration and data cleaning strategies are summarised in Section 6.2. Analyses of the magnetars profile and spectral evolution, the discovery of two kinds of discrete emission mode changing at two different epochs and geometric analyses based on fits to the linearly polarized position angle are presented in Sections 6.3 through 6.5. The implications of our analyses and results are discussed in Section 6.6, with a particular focus on potential physical models that may describe the apparent variations in viewing geometry and polarized emission. We also relate our observations to the transient behaviour of other radio magnetars and high B-field pulsars. A summary of our findings along with concluding remarks are presented in Section 6.7.

6.2 Observations

Following its discovery in March 2020, we began a regular monitoring campaign of Swift J1818.0–1607 with the Parkes UWL receiver system (Hobbs et al., 2020) under the P885 project (PI: F. Camilo). During each observation we typically recorded ~ 10 -minutes of full Stokes pulsar search-mode data covering the full 3328 MHz bandwidth of the UWL with 1 MHz channels and 128 μ s sampling via the MEDUSA backend, where each frequency channel was coherently dedispersed with a dispersion measure (DM) of 706 pc cm⁻³. We created PSRFITS (Hotan et al., 2004) format archives with 1024 phase bins by folding the psrfits-format search-mode data at the topocentric pulse period of the magnetar via DSPSR (van Straten & Bailes, 2011). Calibration and cleaning of the data were performed via the methodology outlined in Lower et al. (2020c). We note for the two observations performed on MJD 58977 and MJD 59009, we used noise diode scans taken 20 minutes after and 30 minutes before the respective Swift J1818.0–1607 observations on these dates. User error prevented the noise diode from activating during the originally scheduled scans. Later observations were not affected by this issue. We tested for inconsistencies in the calibration by measuring the rotation measure (RM) of the polarization spectra at each epoch using both the brute-force method implemented in the RMFIT tool of PSRCHIVE (searched over RM values between -2000 and 2000 rad m⁻² with 4000 steps), and a PYTHON implementation of the direct Stokes Q and U fitting technique described in Bannister et al. (2019). The resulting RM measurements, along with details of each observation are presented in Table 6.1. Note the uncertainties of the RMFIT values are clearly underestimated by about an order of magnitude when compared to those obtained from the Q - U spectral fits. While our recovered RM values deviate from the previously reported value of 1442.0 ± 0.2 rad m⁻² (Lower et al., 2020c), they are consistent with expected variations due to propagation through the ionosphere at the location of Parkes (e.g. Han et al., 2018).

Table 6.1 Parkes UWL observations of Swift J1818.0–1607, along with the number of recorded profile components and rotation measures from RMFIT and direct Stokes $Q-U$ fits.

Observation (UTC)	MJD	Frequency (MHz)	Bandwidth (MHz)	Length (s)	No. profile components	RM (RMFIT) rad m ⁻²	RM ($Q-U$) rad m ⁻²
2020-05-08-18:14:52	58977	2368	3328	639	1	1440.48 ± 0.09	1440.1 ± 0.8
2020-06-09-11:34:36	59009	2368	3328	616	2	1440.86 ± 0.04	1441.7 ± 0.7
2020-07-17-09:35:10	59047	2368	3328	616	2	1441.72 ± 0.05	1439.2 ^{+0.3} _{-0.2}
2020-08-01-13:19:01	59062	2368	3328	2440	2	1441.72 ± 0.05	1447.4 ± 0.4
2020-08-13-10:37:18	59074	2368	3328	617	2	1439.05 ± 0.07	1439 ± 2
2020-08-26-09:06:53	59087	2368	3328	623	3	1439.47 ± 0.04	1440.3 ± 0.4
2020-09-17-04:40:32	59109	2368	3328	1139	2	1443.53 ± 0.04	1440.5 ± 0.4
2020-10-06-07:06:18	59128	2368	3328	618	2	1445.79 ± 0.08	1444.2 ± 0.6

Following this test, we applied the nominal RM of 1442.0 rad m⁻² referenced to the central observing frequency of 2368 MHz to each observation. As our observations are too sparse for a phase connected timing solution to be obtained, all profiles that we show in Figure 6.1 were manually aligned so the total intensity maximum is located at a pulse longitude of 0°.

6.3 Profile and spectral evolution

We show the polarization profiles, linear polarization position angle (PA; Ψ) swings and phase-resolved total intensity spectra for all eight observations in Figure 6.1. The emission profiles on MJD 58978 and 59009 are similar to the profile presented in Lower et al. (2020c) and the subset of those in Champion et al. (2020a) where a single, highly linearly polarized component with a steep spectrum and flat PA were detected. None of the averaged profiles shows evidence for the second component that was occasionally observed by Champion et al. (2020a). However, we later show that a handful of pulses from this previously reported secondary component were detected on MJD 59009. A new profile feature with an inverted spectrum emerged between MJD 59009 and 59047. Reports from other facilities suggest the emission from this profile component is detectable up to frequencies as high as 154 GHz (e.g. Torne et al., 2020b). This inverted-spectrum component persists throughout our later observations. On the other hand, the steep-spectrum component gradually weakens and appears evolve toward more positive values of pulse longitude. By MJD 59128 it is almost completely overlaps with the inverted-spectrum component. Similar longitudinal evolution of individual profile components was

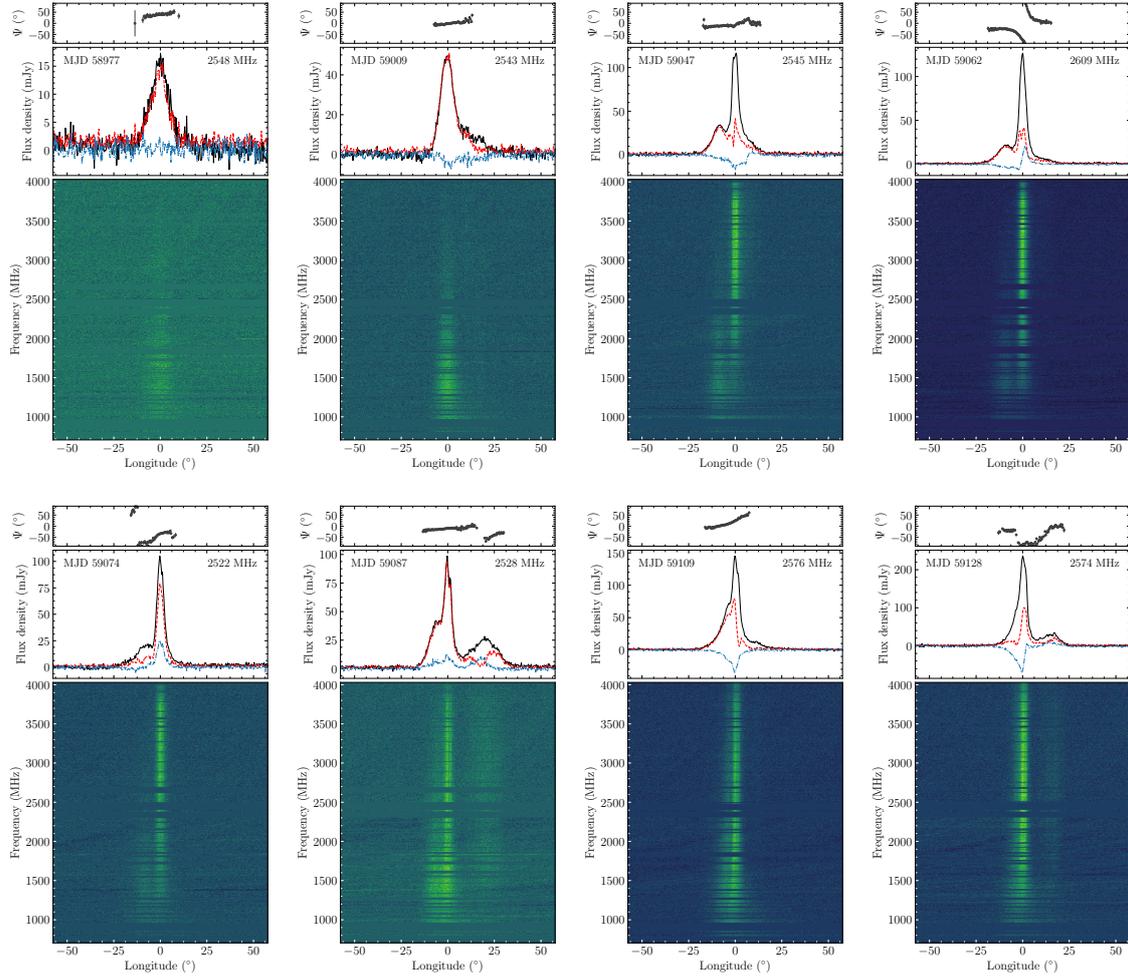


Figure 6.1 Parkes UWL observations of Swift J1818.0–1607. Each plot depicts the linear polarization position angle (Ψ) (top panel), polarization profile with total intensity in black, linear polarization in red and circular polarization in blue (middle panel), and the phase resolved total intensity spectrum (bottom panel).

detected in the pulsed radio emission of XTE J1810–197 following its 2018 outburst (Levin et al., 2019). A third, weaker component that possesses a flat spectrum was detected on MJD 59087 and again as the secondary component on MJD 59128. The 90-degree jump in the PA along with the dip in linear polarization of this component on MJD 59087 are indicative of an orthogonal polarization mode (OPM), as opposed to the PA offset in the secondary component in Figure 7 of Champion et al. (2020a). An OPM is also clearly visible in the leading profile component detected on MJD 59128.

The spectrum of the magnetar has evolved significantly since it was first detected in March 2020. While comparisons of the phase-resolved spectral index would be preferable,

each of the multi-component profiles exhibit variable spectral indices, hence the effects of interstellar scattering would bias our results towards spectra with low-frequency turnovers as the radio flux at low frequencies becomes increasingly spread out as a function pulse longitude. As a result, we were limited to computing the phase-averaged spectral index at each epoch. First, we split the UWL band into thirteen 256 MHz-wide subbands that were then averaged in time and frequency to create a one-dimensional pulse profile for each subband. Corrections to the profile baseline were performed using PSRCHIVE. We then computed the continuum flux density at each subband by averaging over the on-pulse region of each profile as

$$S_\nu = \frac{1}{N_{\text{bin}}} \sum_i^{N_{\text{on}}} S_{\nu,i}, \quad (6.1)$$

where N_{bin} is the total number of phase bins, N_{on} is the number of phase bins covered by the on-pulse region and $S_{\nu,i}$ is the flux at the i -th phase bin. We set the on-pulse window to be between $\phi = -18^\circ$ to 29° for sub-bands above 1.5 GHz, and extend to $\phi = 90^\circ$ below 1.5 GHz in order to account for scatter broadening. For the MJD 59087 and 59128 observations, the extended on-pulse window was used for the full band to accommodate the additional profile components. The flux uncertainty is computed from the normalised root-mean-square (RMS) of the off-pulse region as

$$\sigma_{S,\nu} = \frac{\sqrt{N_{\text{on}}}}{N_{\text{bin}}} \sqrt{\sum_i^{N_{\text{off}}} S_{\nu,i}^2}, \quad (6.2)$$

where N_{off} is the number of bins covering the off-pulse region. We then fit the resulting flux density spectra using either a simple power-law function

$$S(x) = a x^\kappa, \quad (6.3)$$

where a is a scaling parameter, $x = \frac{\nu}{1 \text{ GHz}}$ and κ is the spectral index, or a broken power-law of the form

$$S(x) = a \begin{cases} x^{\kappa_1} & \text{if } \nu \leq \nu_b \\ x^{\kappa_2} x_b^{\kappa_1 - \kappa_2} & \text{otherwise} \end{cases}, \quad (6.4)$$

where $x_b = \frac{\nu_b}{1 \text{ GHz}}$, ν_b is the frequency of the spectral break and κ_1 and κ_2 are the respective spectral indices before and after the spectral break. Posterior distributions for the spectral parameters were sampled using BILBY (Ashton et al., 2019a) as a wrapper for the DYNesty nested sampling algorithm (Speagle, 2020). We assumed a Gaussian likelihood function

Table 6.2 Results from spectral fits and associated log Bayes factors. Observations with only a single spectral index listed are those best described by a simple power-law. Those with two are best fit by a broken power-law.

MJD	$\ln(\mathcal{B}_{\text{SPL}}^{\text{BPL}})$	κ_1	κ_2	ν_b (MHz)
58977	0.7	$-1.7_{-0.3}^{+0.2}$	–	–
59009	0.1	-2.7 ± 0.1	–	–
59047	9.0	-2.0 ± 0.1	0.4 ± 0.2	1801_{-88}^{+111}
59062	8.0	-1.9 ± 0.2	0.4 ± 0.2	1693_{-126}^{+105}
59074	1.5	$-1.2_{-0.3}^{+0.2}$	$0.2_{-0.4}^{+0.5}$	2034_{-418}^{+409}
59087	–0.6	-1.2 ± 0.2	–	–
59109	–0.8	-1.0 ± 0.1	–	–
59128	–1.3	-0.5 ± 0.1	–	–

of the form

$$\mathcal{L}(d|\theta) = \prod_i^N \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{(d_i - \mu_i(\theta))^2}{2\sigma^2} \right], \quad (6.5)$$

where $N = 13$ is the number of frequency subbands, d is the measured flux density, $\mu(\theta)$ is the spectrum model described by parameters θ and $\sigma^2 = \sigma_{S,\nu}^2 + \sigma_Q^2$ is the uncertainty in the flux densities added in quadrature with an additional error parameter (σ_Q) to account for any systematic errors not accounted for in Equation 6.2. We also assumed uniform priors between -10 and 10 for the spectral indices, and a uniform prior spanning 700 MHz to 4000 MHz for the spectral break.

We employed Bayesian model selection to determine which spectral model best described the data. The resulting Bayes factors along with the median recovered values (and associated 68 per cent confidence intervals) for the preferred spectral models are presented in Table 6.2. Our measurements for the single component profiles, in addition to the values of κ_1 on MJD 59047 and 59062, are consistent with the spread of spectral indices between -3.6 and -1.8 presented in [Champion et al. \(2020a\)](#). The recovered values of κ_2 are consistent with the magnetar spectrum being inverted or close to flat at frequencies above 1.6 to 2.0 GHz. Consistently flat spectra were also obtained by observations of Swift J1818.0–1607 by the Deep Space Network between 2.3 and 8.4 GHz on MJD 59045 by ([Majid et al., 2020a](#)), who obtained a spectral index of 0.3 ± 0.2 .

The phase averaged spectrum on MJD 59087 and beyond are best described by the single power-law model, each showing a significant amount of flattening when compared to the previous observations. Although the spectral index of -1.2 ± 0.2 on MJD 59087

is consistent with the value of κ_1 measured on MJD 59074, it appears to have transition back to a single power-law spectrum, albeit one that is much flatter than seen in earlier observations. This can be attributed to a combination of averaging over the additional flat spectrum components detected on MJD 59087 and 59128, and the apparent weakening and increasing level of overlap between the steep- and inverted-spectrum components that we mentioned earlier.

6.4 Emission mode switching

At least two magnetars show evidence for their radio emission switching between multiple, quasi-stable radio profiles (mode-changing) or between an ‘on’ and ‘off’ state (nulling). [Camilo et al. \(2007a\)](#) and [Halpern et al. \(2008\)](#) reported at least two types of discrete state-changes in the single-pulse behaviour of 1E 1547.0–5408, while [Yan et al. \(2018\)](#) noted the Galactic Centre magnetar SGR 1745–2900 would randomly switch between two emission modes in addition to exhibiting nulling. Sudden changes in the profile shape of XTE J1810–197 were also reported by [Camilo et al. \(2007d\)](#) approximately once every 15 hours, while the polarization properties of PSR J1622–4950 could be categorised into four different sub-classes ([Levin et al., 2012](#)). However it is unclear if the phenomena in the latter two magnetars were genuine mode changes or not. Both mode-changing and nulling are thought to be related to same phenomena: variations in (or a complete failure of) the coherent radiation mechanism due to large-scale redistribution of current flows and plasma content in the pulsar magnetosphere ([Kramer et al., 2006a](#); [Wang et al., 2007](#); [Timokhin, 2010](#)). The resulting changes in particle outflows and the associated torque acting to slow the neutron star spin over time have previously been linked to correlated profile shape and spin-down variations in a number of pulsars ([Lyne et al., 2010](#)).

Inspecting the time-phase plot for MJD 59009 in [Figure 6.2](#), it is clear the single-pulse emission of Swift J1818.0–1607 was quasi-periodically switching between a bright mode (B-mode) or a fainter quiet mode (Q-mode). This is distinct from the largely random variations in single pulse flux and jitter often detected in magnetar single pulses (e.g. [Serylak et al., 2009](#); [Maan et al., 2019](#)). Comparing the polarization profiles in [Figure 6.2](#), the B-mode resembles the single component profile presented in [Lower et al. \(2020c\)](#), while the Q-mode is comprised of marginally detected emission at the same longitude as the B-mode and a slightly depolarized bump situated at approximately $+10^\circ$. This bump in the Q-mode profile is positioned at the same pulse longitude as the peak of the inverted-spectrum component detected in later observations and close to the longitude of a depolarized bump seen at high frequencies in [Lower et al. \(2020c\)](#) that was previously dismissed as an arte-

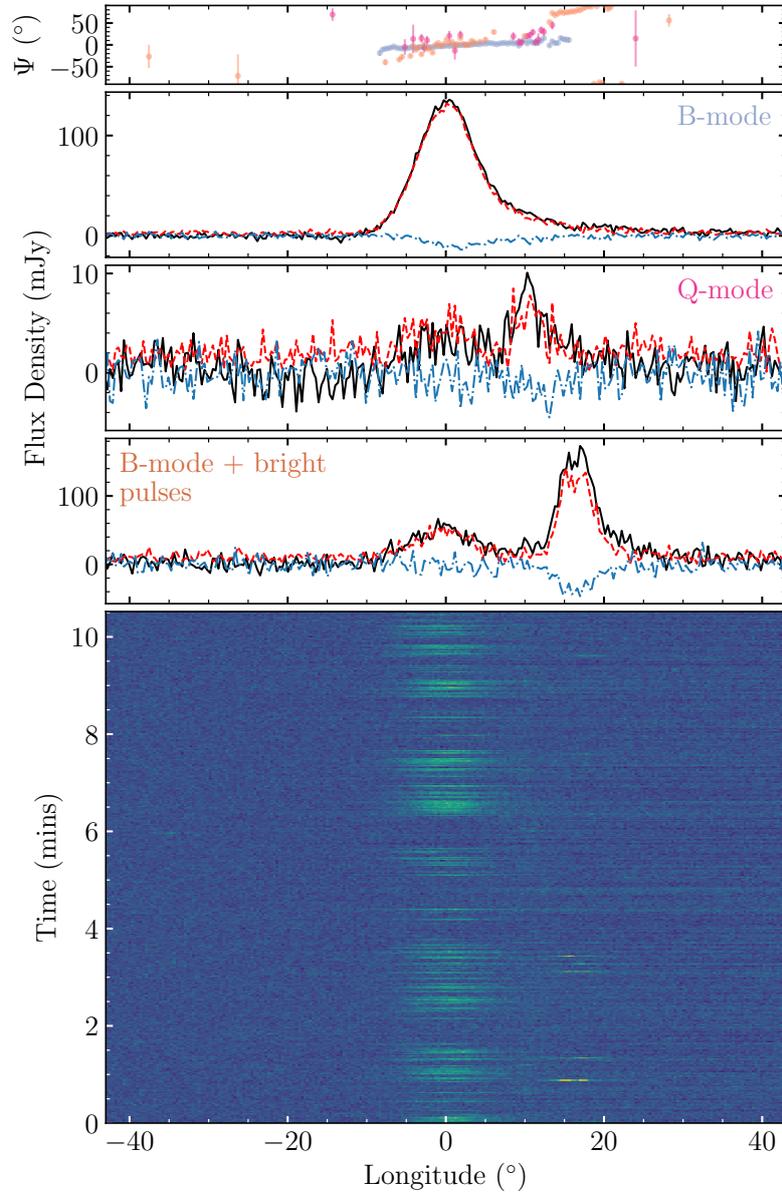


Figure 6.2 Top: comparison of the polarization profiles for the two emission modes detected on MJD 59009 and the profile after averaging only sub-integrations containing bright pulses from the second component. Bottom: stack of single pulses.

fact from residual impulsive RFI. We also detected 10 pulses from a highly intermittent secondary profile component. The fourth panel of Figure 6.2 shows the resulting polarization profile after averaging together the single-pulse archives containing these bright two-component pulses. Both the delay in pulse phase from the primary component and ~ 60 degree offset in the PA swing were seen in a secondary profile component detected

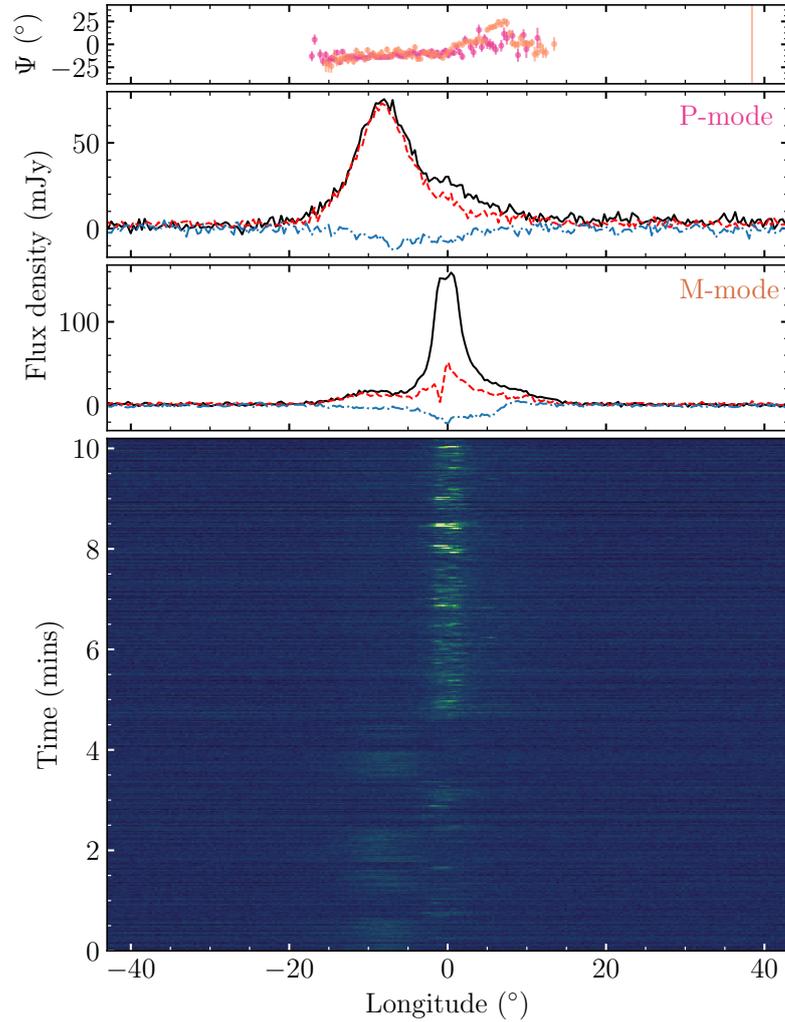


Figure 6.3 Same as Figure 6.2, but for the two emission modes observed on MJD 59047.

by [Champion et al. \(2020a\)](#) around the time of a glitch-like timing event in March 2020, possibly pointing to an increased level of rotational instability around the time of this observation. A plateau in the B-mode profile at the pulse phase the secondary component points to faint pulses from this profile component appearing throughout the observation. In total, we observed 295 rotations spent in the Q-mode and 157 in the B-mode.

We also detected emission mode switching on MJD 59047, however instead of the previous switching between a B- and Q-mode, the time-phase plot shown in Figure 6.3 shows the magnetar varying between two longitudinally distinct modes. We termed these modes the P- and M-modes, as the spectrum of the P-mode resembles the steep spectra often seen in many rotation-powered pulsars while the M-mode exhibits the characteristically flat or inverted spectrum of radio-loud magnetars. Switching between these two modes

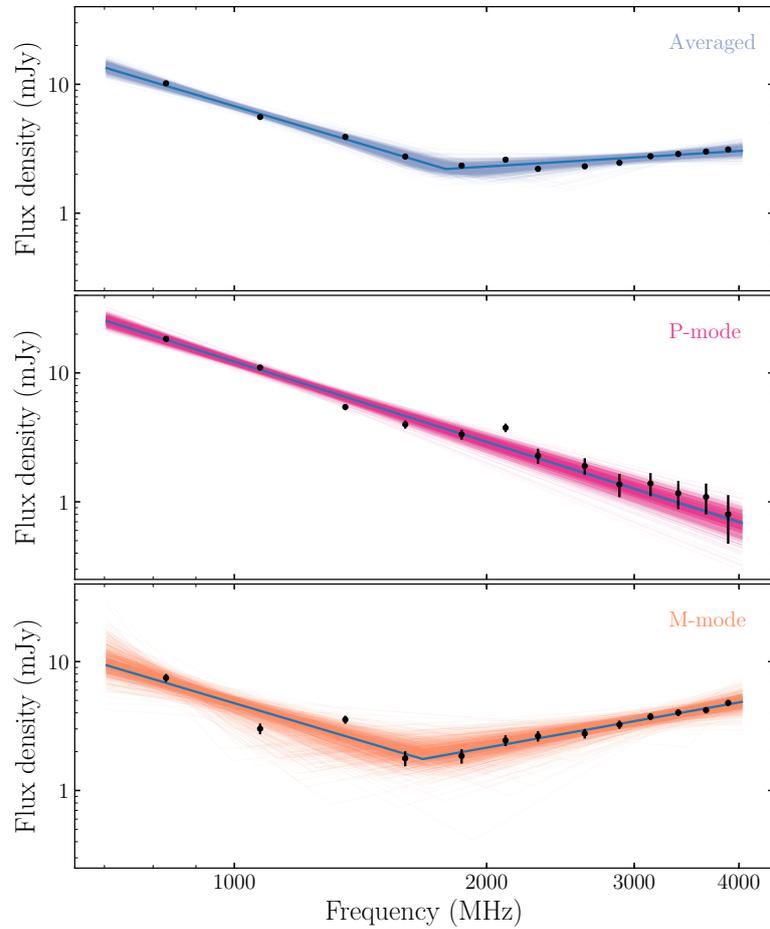


Figure 6.4 Continuum flux densities for the average of the two emission states (top), P-mode only (middle) and M-mode only (bottom). The solid-blue lines indicate the maximum likelihood posterior fit, while the coloured traces represent 1000 random draws from the posterior distributions.

was also detected by [Pearlman et al. \(2020\)](#) who observed Swift J1818.0–1607 with the Deep Space Network two days prior (on MJD 59045) to us. One marked difference to the emission mode switching detected in rotation-powered pulsars, is the mode-changing in Swift J1818.0–1607 was only a temporary phenomenon, as none of our subsequent observations show evidence for discrete switching between modes. Instead, the magnetar appeared to remain in a constant M-mode-like emission state, suggesting whatever mechanism was driving the magnetospheric current variations had stabilised over the course of 15 days.

Using the spectrum fitting techniques outlined in Section 6.3, we found both emission modes detected on MJD 59009 are well described by a simple power-law with steep spectral

indices of $\kappa = -2.6 \pm 0.1$ (B-mode) and $-4.5_{-1.3}^{+1.2}$ (Q-mode). Note, the steeper spectrum of the Q-mode could be an artefact of the low S/N of this emission mode. This is in contrast to the P- and M-modes detected on MJD 59047, where our flux density measurements and spectral fits shown in Figure 6.4 clearly show the P-mode has a steep spectrum with $\kappa = -2.0 \pm 0.1$ and the M-mode shows evidence of a spectral break at ~ 1652 MHz and transition to an inverted spectrum, with corresponding spectral indices of $\kappa_1 = -2.0_{-0.6}^{+0.4}$ and $\kappa_2 = 1.2 \pm 0.3$ before and after the break. The consistency between the P-mode spectral index and M-mode pre-spectral break index indicates the magnetar continues to emit weak radio pulses from the P-mode-component while the M-mode is dominant.

Lastly, we checked for differences in the RM between the various emission modes that can arise from a variety of physical processes, such as the superposition of OPMs with different spectral indices and propagation effects within the neutron star magnetosphere (Noutsos et al., 2009; Ilie et al., 2019). While we do successfully recover a RM of $= 1441.0 \pm 0.6 \text{ rad m}^{-2}$ for the B-mode detected on MJD 59009, the RM was unconstrained for the Q-mode owing to the low level of emission associated with this mode. For the P- and M-modes detected on MJD 59047, we obtained respective RM values of $1441.6 \pm 0.7 \text{ rad m}^{-2}$ and $1440.2 \pm 0.9 \text{ rad m}^{-2}$. There is a significant amount of overlap between the posteriors for these two modes at the 68 per cent confidence interval, suggesting any propagation effects between the the two modes are negligible.

6.5 Polarization properties and geometry

Earlier works noted Swift J1818.0–1607 possessed a relatively flat PA, potentially pointing to our line-of-sight only grazing the emission cone edge (Lower et al., 2020c; Champion et al., 2020a). Similarly flat PAs were also detected across the first three observations shown in Figure 6.1, however the PA swings across the last five epochs each differ dramatically, bearing a striking resemblance to the S-shaped swing expected from the simple rotating vector model (RVM) of Radhakrishnan & Cooke (1969). Under the RVM, the sweep of the PA is a purely geometric effect caused by the changing angle between the projected dipole magnetic-field direction and our line of sight. It can be expressed in terms of the magnetic inclination angle (α ; the angle between the spin and magnetic axes) and the angle between the spin axis and our line of sight (ζ) as

$$\tan(\Psi - \Psi_0) = \frac{\sin \alpha \sin(\phi - \phi_0)}{\sin \zeta \cos \alpha - \cos \zeta \sin \alpha \cos(\phi - \phi_0)}, \quad (6.6)$$

Table 6.3 RVM-fits to the data. MJDs with a † include corrections for orthogonal polarization modes.

MJD	α ($^{\circ}$)	β ($^{\circ}$)	ϕ_0 ($^{\circ}$)	Ψ_0 ($^{\circ}$)	ζ ($^{\circ}$)
58977	106^{+29}_{-36}	-67^{+30}_{-41}	-3^{+35}_{-32}	15^{+37}_{-35}	24^{+44}_{-18}
59009	93^{+39}_{-31}	-37^{+14}_{-20}	20^{+14}_{-19}	25^{+21}_{-26}	49^{+58}_{-36}
59047	106^{+23}_{-29}	-71^{+30}_{-32}	4^{+33}_{-42}	2^{+36}_{-42}	20^{+45}_{-16}
59062	82 ± 9.4	$2.80^{+0.08}_{-0.13}$	1.69 ± 0.08	$77.2^{+0.9}_{-0.8}$	85 ± 10
59074	115^{+31}_{-46}	-9^{+4}_{-2}	-7 ± 1	-71^{+4}_{-6}	104^{+37}_{-46}
59087†	113^{+11}_{-9}	$-11.2^{+0.9}_{-0.6}$	$25.3^{+0.7}_{-0.8}$	-40 ± 3	102^{+12}_{-9}
59109	144^{+14}_{-22}	-6 ± 3	$4.7^{+0.5}_{-0.3}$	-48^{+3}_{-2}	138^{+19}_{-25}
59128†	71 ± 28	$-8.5^{+2.0}_{-0.8}$	11.1 ± 0.8	-40^{+5}_{-3}	62^{+28}_{-26}

where ϕ_0 is the pulse longitude at which $\Psi = \Psi_0$, i.e, the PA of the pulsar spin axis projected onto the plane of the sky. The difference between ζ and α is the angle of closest approach between our line of sight and the magnetic axis ($\beta = \zeta - \alpha$), hereafter referred to as the magnetic impact angle. While the RVM is only truly valid in the case of an unchanging, axisymmetric dipole magnetic field, the geometric interpretation of the model can potentially provide some insight to the processes driving the PA variations (Everett & Weisberg, 2001; Johnston & Kramer, 2019).

6.5.1 Viewing geometry

Using the RVM and a Gaussian likelihood function, we fitted each of the PA swings shown in Figure 6.1. We assumed uniform priors on all RVM parameters, except for ϕ_0 where we employed a Gaussian prior centred at 0° with a width of 45° . This constrained prior allows us to avoid the ambiguity in which magnetic pole the polarized radio emission originates, as we do not know the sense of the magnetar’s rotation. For the observations on MJD 59062 and 59074 we applied a $+180^{\circ}$ and -180° phase jump respectively to PA values below -6° in order to have a smooth PA swing across the pulse profile. We also corrected the 90° jump in the PA swings on MJD 59087 and 59128 due to OPM transitions by subtracting -90° from the measured PA values at $\phi \leq 9^{\circ}$. The results of our RVM-fits are presented in Table 6.3. In Figure 6.5 we show the PA swings from the last five epochs; our overlaid RVM-fits are in excellent agreement with the data.

Our best constrained values of the geometry from the MJD 59062 ($\alpha = 82^{\circ}$, $\beta = 3^{\circ}$) and 59087 ($\alpha = 112^{\circ}$, $\beta = -11^{\circ}$) observations are highly inconsistent, as could already

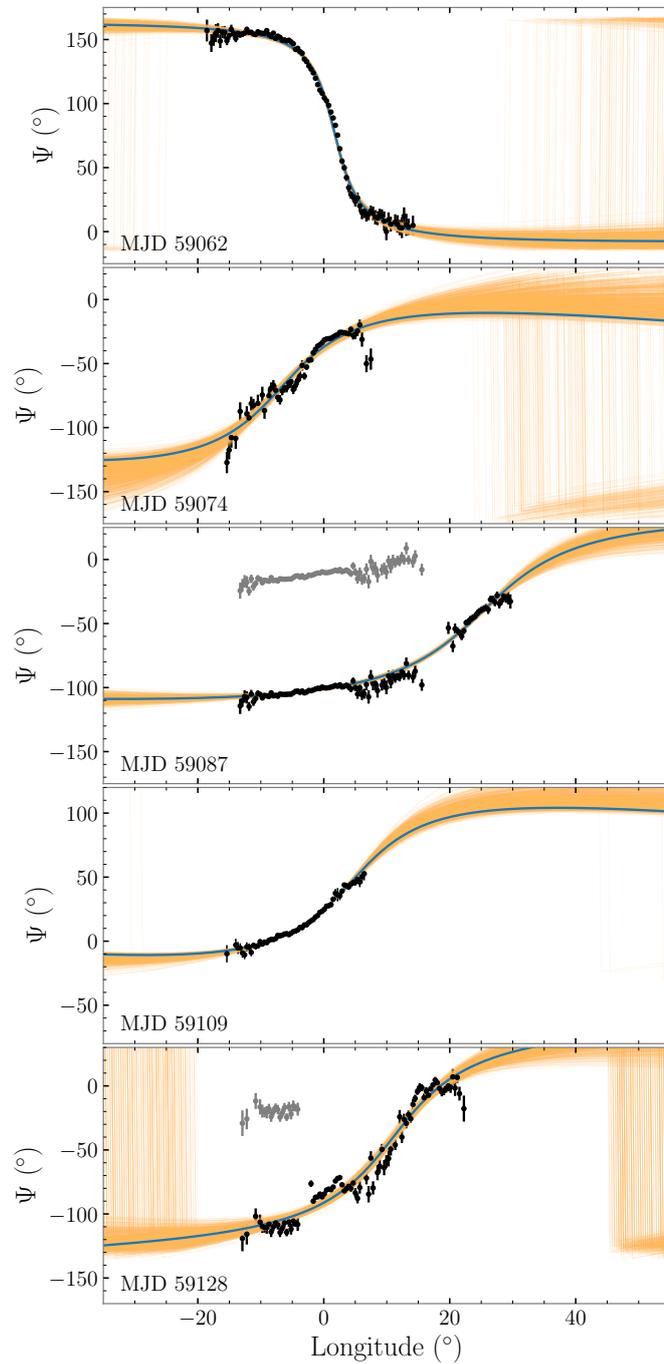


Figure 6.5 RVM fits to the five RVM-like PA swings (black points) with the maximum likelihood a posteriori fit (blue) along with traces generated from 1000 random draws from the posterior distributions (orange). Grey points indicate PA values prior to adding an OPM correction.

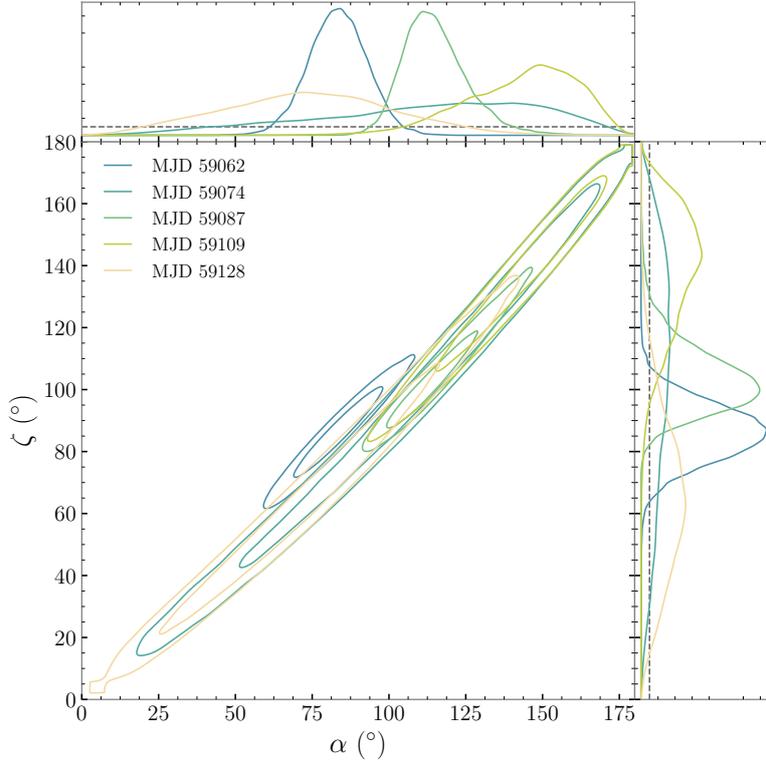


Figure 6.6 One- and two-dimensional posterior distributions of α and ζ . Contours indicate the 68% and 95% confidence regions. Grey-dashed lines in the one-dimensional posteriors indicate the priors.

be discerned from the opposite sweep of the PAs on these dates. The relatively flat PA swings and narrow pulse duty-cycles seen on MJD 58977, 59009 and 59047 resulted in our recovered values for α and ζ being relatively unconstrained although the positive PA gradients indicate $\beta < 0$ at these epochs. Figure 6.6 and 6.7 shows the one- and two-dimensional posterior distributions of α , ζ and β from our fits to the PAs in Figure 6.5. It is clear that both α and ζ remain largely consistent between the four observations that show positive PA gradients. As the marginalised ζ posteriors for MJD 59062 and 59087 share a significant amount of overlap at the 68% confidence level and the α posteriors do not, the most likely explanation for the flipped PA swing and inferred β on MJD 59062 is a sudden change in α that occurred between MJD 59047 and MJD 59062, that subsequently reversed sometime prior to MJD 59074. We discuss the implications and describe probable causes of this effect below in Section 6.6. If we ignore the results from MJD 59062, then we can combine the α and ζ posteriors at every other epoch to obtain improved measurements of $\alpha = 112^\circ_{-9}^{+7}$ and $\zeta = 99^\circ_{-10}^{+7}$, which in turn correspond to $\beta = -12.9^\circ_{-0.7}^{+0.6}$.

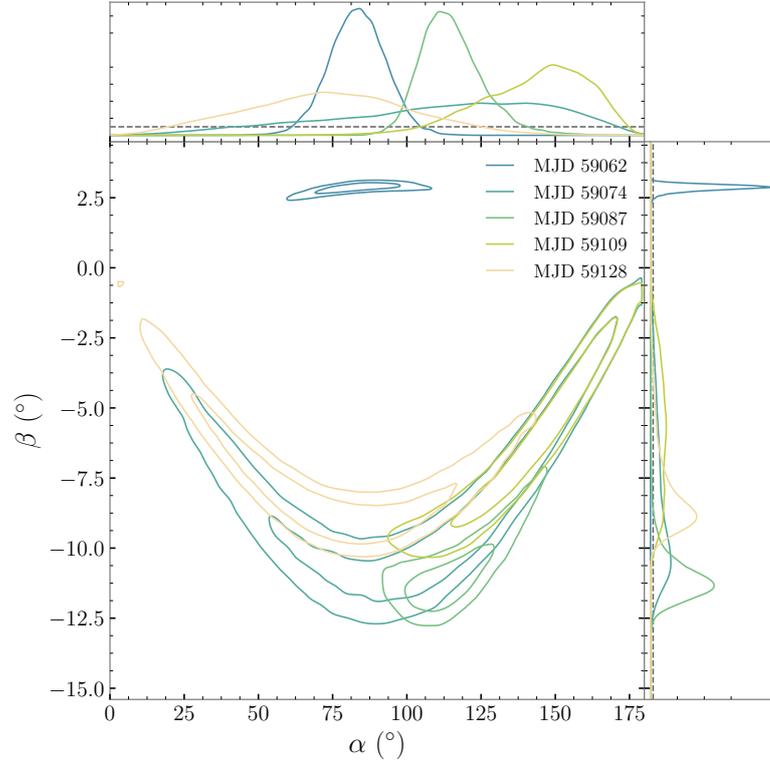


Figure 6.7 One- and two-dimensional posterior distributions of α and β .

Given the inferred geometry, the radio pulses must originate from high above the neutron star surface in order to explain the average profile width (at the 10% flux level) of $W_{10,\text{avg}} = 34.1^\circ$. Using the measured values of α and ζ , we can infer a minimum geometric emission height by first computing the half-opening angle of the emission cone (ρ ; Gil et al. 1984)

$$\cos \rho = \cos \alpha \cos \zeta + \sin \alpha \sin \zeta \cos(W/2), \quad (6.7)$$

where W is the pulse width, taken to be W_{10} (in units of rad) in our case. Assuming the emission extends to the last open field line and a fully active polar cap with symmetric emission about the pole, the emission height, h_{em} , can be derived via (Rankin, 1990)

$$\rho = 3\sqrt{\frac{\pi h_{\text{em}}}{2Pc}}, \quad (6.8)$$

where $P = 1.3635$ s is the spin period of the magnetar and c is the vacuum speed of light. From the inferred geometry and W_{10} we obtain a minimum emission height of 3800 km, i.e close to 6 per cent the light-cylinder radius of Swift J1818.0–1607 ($r_{\text{lc}} = 6.5 \times 10^4$ km).

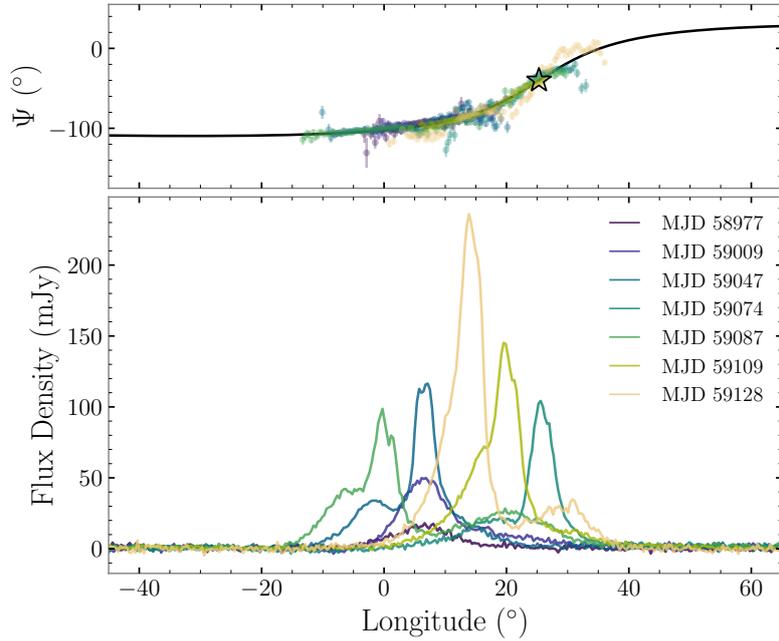


Figure 6.8 Comparison of the PA swings (top) and total intensity profiles (bottom) of the PA aligned profiles. The black coloured line and star in the top panel represents the median RVM-fit and corresponding position of (ϕ_0, Ψ_0) for the PA swing on MJD 59087.

6.5.2 Position angle alignment and emission heights

Visually, the PA swings that are shown for MJD 59074 onward in Figure 6.5 appear similar, and could easily be aligned by the addition of 90° jumps in PA and shifts in pulse longitude. A similar argument could be made for all of the PA swings prior to MJD 59062 since they all show evidence of shallow positive PA gradients. To investigate whether it is possible to align the PAs, we first visually aligned each PA swing by adding a -135° jump to MJD 58977 and -90° jumps to MJD 59009, 59047 and 59109 respectively. We then measured the longitude offsets required to align the PAs by performing a least-squares fit to the data assuming values of $(\alpha, \zeta) = (112^\circ, 99^\circ)$ from the combined posteriors, and $(\phi_0, \Psi_0) = (25.5^\circ, -40.2^\circ)$ from the RVM fit to MJD 59087. The PAs and total intensity profiles after applying the resulting phase offsets are displayed in Figure 6.8. There are two possible ways to interpret the PA aligned profiles: longitudinal wandering or oscillating of the emission patch over time, or temporal evolution of the emission height. Longitudinal motion of the emission patch would imply the the magnetar possesses a largely unfilled emission cone with an opening angle that is difficult to reconcile with our inferred magnetic geometry and its rotation period of ~ 1.4 s. Hence, for the remainder of this section we focus on the more likely possibility of a changing emission height.

Table 6.4 Longitude offsets and relative emission heights.

MJD	$\delta\phi$ ($^{\circ}$)	$h_{\text{em}} - h_{\text{em,MJD 59074}}$ (km)	$(h_{\text{em}} - h_{\text{em,MJD 59074}})/r_{\text{lc}}$
58977	18.6	5200	0.08
59009	19.1	5400	0.08
59047	18.8	5300	0.08
59074	0	0	0
59087	25.5	7200	0.11
59109	5.7	1600	0.02
59128	11.7	3300	0.05

Pulsar emission theory predicts radiation produced nearer to the magnetic pole will originate from magnetic field lines closer to the neutron star surface (see, e.g. [Yuen & Melrose, 2014](#)). If the emission region is symmetric about the magnetic meridian, then relativistic aberration and retardation effects will cause the observed PA to lag behind the total intensity profile ([Blaskiewicz et al., 1991](#)). The emission height relative to the centre of the neutron star (h_{em}) can be inferred from the magnitude of this delay expressed in terms of pulse longitude ($\delta\phi$, in units of rad) and radius of the light cylinder, r_{lc} , as

$$h_{\text{em}} = \frac{r_{\text{lc}}}{4} \delta\phi = \frac{Pc}{8\pi} \delta\phi. \quad (6.9)$$

Figure 6.8 shows that the observation on MJD 59074 must have the lowest emission height, as the pulse profile is almost aligned with the inflexion point of the RVM fit. We therefore take this observation as a reference for computing relative emission heights noting that the absolute height is difficult to ascertain as the location on the profile of the pole crossing is unclear. Table 6.4 shows the longitude offset between the profiles relative to the observation on MJD 59074 and hence the inferred relative values of h_{em} expressed in both km and as a fraction of the light cylinder radius. We choose not to include uncertainties because the main point is to demonstrate indicative changes in emission height. The table shows there is substantial variation in emission height between the epochs and no particular trend with time. Why the emission height should change in this way is unclear, but the fact that the profile components persist implies that the same field lines are being illuminated over the range of heights.

It is evident the polarization profiles corresponding to epochs with smaller inferred emission heights possess the largest variations in polarization fraction. For instance, the

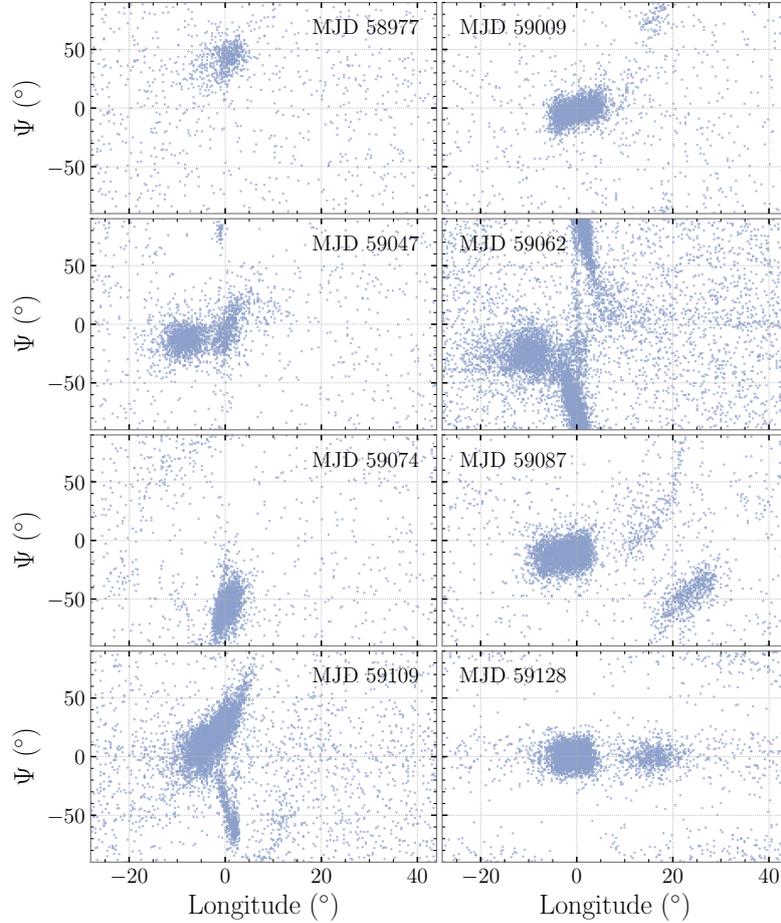


Figure 6.9 Pulse and phase resolved polarization position angle (Ψ).

leading, steep-spectrum component normally has close to 100 percent linear polarization, however on MJD 59074 – the observation with the smallest inferred emission height – its polarization fraction is less than half of what it is at every other epoch. This is consistent with observations of rotation-powered pulsars, where radio emission emitted lower in the magnetosphere is more likely to be subject to a larger amount of magnetospheric propagation effects (e.g. [Smith et al., 2013](#)).

6.5.3 Polarization modes

Many normal pulsars emit linearly polarized radio waves in two modes that are usually orthogonal to one another ([Backer et al., 1976](#)). It is believed these modes arise from propagation effects within the neutron star magnetosphere, such as refraction and birefringence ([McKinnon, 1997](#); [Petrova, 2001](#)). If two or more OPMs exist, then the process

of averaging over many rotations can suppress the observed linear polarization in pulsar profiles (e.g. Karastergiou et al., 2002). This could explain the apparent depolarization of the inverted spectrum profile component detected on MJD 59047 onward. We tested this idea by studying the distribution of PA values at each phase bin across the pulse profile. To minimise spurious contributions from noise and unaccounted RFI, we imposed a threshold where the linear polarization of a given phase bin must be a factor of 2.5 times greater than the off-pulse RMS when calculating the PA. Scatter plots showing the PA distributions for all eight epochs are presented in Figure 6.9.

In general, the scatter plots largely follow the PA swings depicted in Figure 6.1. This is not surprising for the profiles/profile-components that show a high amount of linear polarization, as the presence of OPMs would result in depolarization. Aside from the known OPM on MJD 59087, we find evidence of additional OPMs on MJD 59047 at longitudes between -5° to 0° , as indicated by the small cluster of points that have an approximately $+90^\circ$ offset in PA from the majority of the scatter plot, and a possible OPM on around longitudes close to 0° on MJD 59062. The offset clump of PA values at longitudes between 10 to 20° on MJD 59009 can be attributed to the handful of bright pulses detected from the secondary profile component shown in the fourth panel of Figure 6.2. In general, there is a notable lack of additional polarization modes in the components that have low linear polarization fractions. This could be due to pulses displaying emission from additional modes being intermittent and we simply did not catch a large amount of these pulses in our relatively short observations. Longer observations performed by other facilities may be able to place stronger constraints on the presence of any additional polarization modes.

While the majority of the PA distribution on MJD 59109 follows the expected curve seen in the average PA, there are a number of points that follow a branching PA swing that bends away from the bulk distribution. Intriguingly, the slope of the branch appears to match that of the PA swing (and overall PA distribution) observed on MJD 59062. Similar branching behaviour has been seen in the PA distributions of some rotation-powered pulsars (e.g. Figure 4 of Ilie et al., 2020). Remarkably, after visually aligning of the two PA distributions by adding a $+3^\circ$ offset to MJD 59062, the slope of the drifting branch matches the downward portion of the PA distribution of MJD 59062. This naturally raises the question: did we observe sporadic pulses from the same, reversed PA emission mode that was detected on MJD 59062?

Inspection of individual pulses associated with this ‘abnormal’ polarization mode reveals the majority exhibit PAs with a continuous downward drift as a function of pulse longitude, while a handful show evidence of an initially upward drifting PA followed by

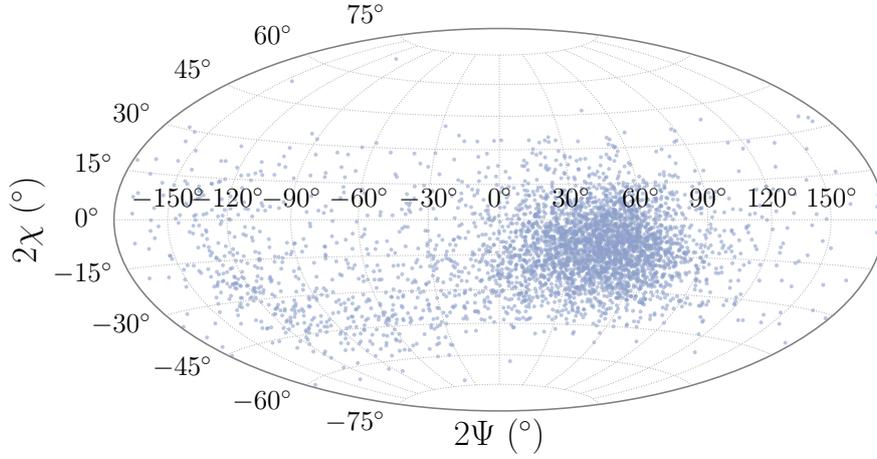


Figure 6.10 Hammer equal-area projection of the polarization position angle (Ψ) and ellipticity angle (χ) distributions from MJD 59109 on the Poincaré sphere.

an apparent OPM jump to the tail of the downward-drifting PA distribution. In general, they all show a lower linear polarization fraction compared to pulses from the ‘normal’ mode and significantly increased amounts of circular polarization. Dyks (2020) devised a phenomenological model in which similar behaviour can originate from the passage of the emission patch along a great circle close to one of the Stokes V poles when projected onto the Poincaré sphere. We tested whether such a passage is present within our data by visually inspecting the position angle and corresponding ellipticity angle distributions plotted on the Poincaré sphere in Figure 6.10. Most of the polarization distribution is concentrated in a blob centred near $(2\Psi, 2\chi) = (45^\circ, -15^\circ)$, however the low-density distribution with negative values of Ψ , i.e. values associated with the downward drifting branch, appears to trace out a rough circular pattern similar to those presented in Figures 2 and 3 of Dyks (2020). This suggests the pulses associated with the downward-drifting PA branch are not associated with the reversed PA swing detected on MJD 59062. Instead they may represent a sample of pulses that experienced a propagation effect within the magnetosphere that masquerades as a smeared OPM in the 2-dimensional PA-longitude plot.

6.6 Discussion

Radio-loud magnetars are unusual in that their flat or inverted spectra means they are detectable as pulsars at millimetre-wavelengths (e.g. Camilo et al., 2007d). Hence it was surprising when Swift J1818.0–1607 was found to possess a steep, negative spectral index. However, given its similarities to the population of high B-field pulsars, we speculated in

Lower et al. (2020c) that the current radio outburst may progress in a similar fashion to the 2016 magnetar-like outburst of PSR J1119–6127 (Majid et al., 2017), and the spectrum could begin to flatten over the months following its discovery. An earlier, multi-wavelength observation by the Deep Space Network (MJD 58947) reported a possible flattening of the spectrum (Majid et al., 2020c), and the apparent trend toward smaller spectral indices over time found by Champion et al. (2020a) seemingly pointed to the spectrum following this prediction. However, our first two spectral index measurements listed in Table 6.2 appear to be in conflict with this hypothesis, and it was only after the emergence of a new profile component bearing an inverted spectrum that the phase-averaged spectral index showed any sign of flattening. Our measured spectral index for this component ranges between $\kappa = -0.2$ to $+0.7$, similar to the those of other radio-loud magnetars (Levin et al., 2012; Dai et al., 2019), and enabled pulses from Swift J1818.0–1607 to be detected up to millimetre wavelengths (Torne et al., 2020b). Intriguingly the negative reported spectral index at these high wavelengths, combined with flatter, but still negative spectral indices measured between 6 and 39 GHz by Effelsberg and the Deep Space Network (Liu et al., 2020; Pearlman et al., 2020), indicate Swift J1818.0–1607 possesses a high-frequency spectral turnover. Both SGR 1745–2900 and XTE J1810–197 were detected at similarly high radio frequencies following their 2013 and 2003/2018 outbursts (Torne et al., 2015; Pennucci et al., 2015; Camilo et al., 2007d; Torne et al., 2020a) and also showed evidence of similar spectral behaviour, indicating high-frequency turnovers may be a common feature of the magnetar radio emission mechanism.

In addition to developing a flat-spectrum component, we also detected two distinct types of emission mode switching at two separate epochs, along with dramatic variations in the position angle swing. However, Swift J1818.0–1607 is not the only magnetar found to exhibit mode switching. The mode changes in the single pulses from SGR 1745–2900 are a subtle effect, manifesting as slight changes in the leading edge of its profile (see Figure 3 of Yan et al., 2018). In contrast, 1E 1547.0–5408 has been seen to undergo at least two types of transient profile events: bright bursts followed by emission appearing at slightly earlier pulse longitudes before ‘recovering’ back to its initial position (Figure 2 of Camilo et al., 2007a), and discrete switching to and from an emission mode where the profile grows an extra hump on its trailing shoulder (see Figure 5 of Halpern et al., 2008). The latter mode appears somewhat similar to the P- and M-mode switching we detected on MJD 59047, however the lack of spectral analyses of the 1E 1547.0–5408 precludes a more direct comparison. We can however draw some parallels between the modes of Swift J1818.0–1607 and the curious behaviour of the high B-field pulsar PSR J1119–6127, where

a number of one-off profile variations were observed by [Weltevrede et al. \(2011\)](#) following a large glitch in 2007. This included a transient secondary profile component that lags the primary by $\sim 30^\circ$, bearing a somewhat similar profile shape and polarization fraction to the secondary component we detected in Swift J1818.0–1607 on MJD 59009. They also detected highly sporadic pulses similar to those from rotating radio transients (RRATs) during two separate epochs where the pulsar was observed at two different frequency bands: 4 bright pulses during a 20-cm observation, and a handful at 10-cm. It was argued the rate at 10-cm must be much higher than at 20-cm as the pulsar was rarely observed at this frequency band. This apparent increased detection rate could be interpreted as the RRAT-like pulses possessing a more magnetar-like, inverted spectral index. Intriguingly, no reported enhancement to the pulsars X-ray emission was associated with this glitch ([Gögüş et al., 2016](#)), unlike the 2016 glitch that was associated with a magnetar-like outburst ([Archibald et al., 2016](#)). [Dai et al. \(2018\)](#) found PSR J1119–6127 exhibited dramatic variations in its polarisation properties during the 2016 outburst, in particular the transient secondary component, showed a similar amount of polarization variations as Swift J1818.0–1607. One notable difference between the polarization variations in Swift J1818.0–1607 and PSR J1119–6127 is the latter showed extreme deviations from its normally flat PA swing, exhibiting a variety of non-RVM-like variations over the course of a few days. Similar strong variations in the polarization fraction and PA swing of XTE J1810–197 were observed after its 2018 outburst that again deviate significantly from the predictions of a simple RVM model ([Dai et al., 2019](#)).

If these variations in the pulse profile are associated with fluctuations in the magnetospheric currents, then we might expect there to be some correlation with the spin-down behaviour or high-energy activity of the magnetar. To place the emergent phenomena in context, we have plotted the inferred change in spin-frequency measured at each epoch (referenced to the spin-frequency on MJD 58977) after subtracting off a constant spin-down rate of $-2.37 \times 10^{-12} \text{ s}^{-2}$ in Figure 6.11. Also shown are epochs where high-energy bursts were detected by the Burst Alert Telescope on board *Swift*. Visually, it is evident that at least two variations in the spin-down rate have occurred over the timespan covered by our observations, as indicated by the relatively sharp changes in $\Delta\nu$. The upward trend between MJD 59047 and 59087 could be a result of decreased particle outflows following the stabilisation of the inverted-spectrum component (and the associated magnetospheric currents; see e.g. [Kramer et al., 2006a](#)), while the flattening from MJD 59087 to 59128 could be associated with the evolution and eventual overlapping of the steep-spectrum component into the inverted-spectrum component. Alternatively, these spin-frequency

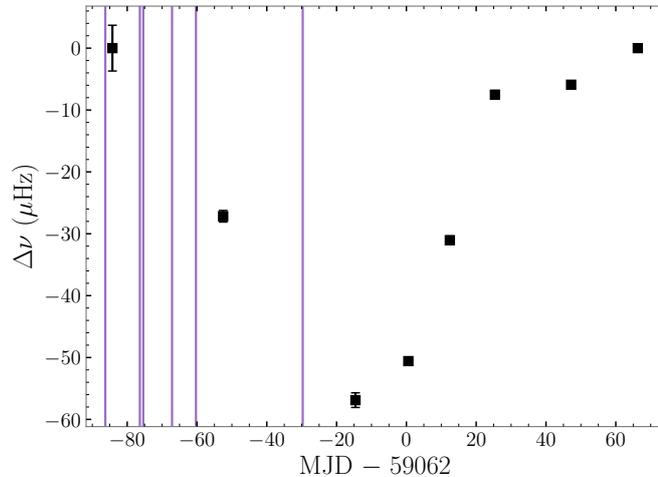


Figure 6.11 Variations in the spin-frequency of Swift J1818.0–1607 over time. Vertical lines correspond to high-energy bursts detected by *Swift* (Barthelmy et al., 2020; Gronwall et al., 2020; Bernardini et al., 2020).

variations could be a result of the magnetar changing spin-down modes similar to what was reported by Champion et al. (2020a). The GCN describing the hard X-ray/gamma-ray burst on MJD 59032 reported a peak count rate of ~ 3000 counts/s $^{-1}$, twice that of the initial burst that led to the discovery of Swift J1818.0–1607 (Evans et al., 2020). It’s possible the resulting magnetic field reconfiguration associated with this burst triggered the emergence of the inverted spectrum profile component, initially through the transient P- and M-mode switching that we detected on MJD 59047. Given our relatively sparse observing cadence, we cannot confirm a causal relationship between these two events. Facilities with high observation cadences may be able to confirm or rule out a potential association.

Measurements of the magnetic geometries of magnetars is useful for both comparing predictions of how their magnetic fields may evolve over long timescales (Tauris & Manchester, 1998; Viganò et al., 2013; Gourgouliatos & Cumming, 2014), and for understanding their outburst mechanism (Perna & Pons, 2011; Rea et al., 2012; Li et al., 2016). However, only a handful of magnetars have had their magnetic geometries constrained through radio polarimetry and fitting of their X-ray profiles, and various arguments have been made (Kramer et al., 2007; Camilo et al., 2008; Levin et al., 2012). From our geometric fits to the PA of Swift J1818.0–1607, we inferred a magnetic and viewing geometry of $(\alpha, \zeta) = (112^\circ_{-9}^{+7}, 99^\circ_{-10}^{+7})$, indicating it is an orthogonal rotator. A similar geometry was also inferred from polarimetry of the prototypical radio-magnetar XTE

J1810–197. Camilo et al. (2007d) found both nearly aligned ($\alpha \sim 4^\circ$ and $\beta \sim 4^\circ$) and close-to-orthogonal ($\alpha \sim 70^\circ$ and $\beta \sim 20\text{--}25^\circ$) RVM-fits were both consistent with the data, depending on whether or not an OPM jump was included for the PA swing across the inter-pulse. However, Kramer et al. (2007) argued a single RVM was insufficient to simultaneously fit both the main and interpulse. Instead, they found that two separate fits to the individual components returned a consistent $\zeta = 83^\circ$ despite having recovering different values of α and β for the main-pulse ($\alpha \sim 44^\circ$, $\beta \sim 39^\circ$) and interpulse ($\alpha \sim 77^\circ$, $\beta \sim 6^\circ$). Perna & Gotthelf (2008) and Bernardini et al. (2011) obtained similar constraints on the angles the line-of-sight and X-ray hotspot pole make with the spin-axis when the magnetar was in its outburst and quiescent states. The deviation of both XTE J1810–197 and Swift J1818.0–1607 from being aligned rotators adds further credence to the argument that their broad radio profiles must be due to emission originating at large heights within the magnetosphere. It also rules out the rapid magnetic and spin axes alignment hypothesis we put forward in Lower et al. (2020c) as a possible explanation for the apparent young age of Swift J1818.0–1607 despite the lack of an obvious associated supernova remnant.

A complication to our geometric interpretation is the flipped PA swing direction (negative gradient instead of positive) we detected on MJD 59062. Naively we could interpret this phenomena as either radio emission originating from the antipodal magnetic pole or our line of sight having undergone a latitudinal crossing of the magnetic pole. Under the RVM, emission from the antipodal pole of the neutron star would exhibit a PA swing with the opposite sign, something that has been observed in a handful of pulsars where emission from both poles are detected as a main pulse and an interpulse (e.g. Johnston & Kramer, 2019). Similarly, geodetic precession of the relativistic binary pulsar PSR J1906+0746 resulted in a sign flip of its PA swing as the magnetic pole crossed our line of sight (Desvignes et al., 2019). While a flipping of the emission to the opposite magnetic pole of Swift J1818.0–1607 could in principle explain the flipped PA swing, the averaged total intensity profile and spectra remains almost identical to that seen during the previous observation, making this scenario unlikely as the magnetic field and current configurations would have to be identical at both polar caps. If the shape of Swift J1818.0–1607 deviates from spherical symmetry due to crustal or magnetic stresses the spin axis can become offset from the total angular momentum vector. This would cause the spin-axis to undergo free precession about the total angular momentum vector, resulting in both an apparent latitudinal and longitudinal evolution of the magnetic axis over time (Pines, 1974). However, free precession also presents an unlikely explanation for the profile and

geometric variations. The short precession timescale required to explain our data would introduce periodic spin-down variations that are not detected in the high cadence timing by [Champion et al. \(2020a\)](#) and [Hu et al. \(2020b\)](#). Also, if the first timing event reported by those two studies is a true spin-up glitch, then free-precession is even more unlikely as the presence of pinned vortices within the neutron star core would rapidly dampen any precession ([Shaham, 1977](#)). Further weight against the emission flipping between poles and the free precession arguments comes from both assuming a static, unchanging magnetosphere, where the observed profile variations are purely due to changes in the viewing geometry, whereas we have shown the magnetic and viewing geometries remain largely unchanged. Additionally, our assumption that Swift J1818.0–1607 has a predominately dipole magnetic field geometry may be incorrect.

There are numerous theoretical and observational studies throughout the literature that point to magnetars possessing dynamic magnetic fields, where non-axisymmetric field geometries, higher-order multipoles and closed magnetic loops are suggested play an important role in describing the observed phenomenology ([Thompson & Duncan, 1993](#); [Thompson et al., 2002](#); [Beloborodov, 2009](#)). While complex multipole fields are likely to be present close to the surfaces of most neutron stars, the success of the RVM in describing the PA swings we observe suggests a more simplistic field geometry is associated with the radio emitting region of Swift J1818.0–1607. NICER observations of Swift J1818.0–1607 by [Hu et al. \(2020b\)](#) showed the X-ray profile exhibits an unusually high pulse fraction for a profile with only a single component and noted it would be difficult to reproduce with the canonical two antipodal hotspot model. They suggested this may instead be evidence the pulsed X-ray emission originates from either a single distorted surface hotspot or a two-component hotspot with differing temperatures. A possible framework for describing such a hotspot configuration is provided by the magnetar corona model of [Beloborodov & Thompson \(2007\)](#), where the high-energy and radio emission originates from either the closed magnetic field loops or open field lines emerging from two sites (starspots) on the neutron star surface – somewhat analogous to coronal loops in the solar magnetic field that link pairs of sunspots. Assuming this coronal loop interpretation holds true for Swift J1818.0–1607, we can explain the flipping of the position angle swing detected on MJD 59062 as being due to highly intermittent switching of the emission region between a more active ‘primary’ and less active ‘secondary’ starspot. A similar hypothesis was put forward by [Kramer et al. \(2007\)](#) to explain the PA swing of XTE J1810–197 during its 2003 outburst, where their preferred, dual RVM-fits were speculated to be evidence of radio emission originating from two active poles within a global multipolar field. Interpret-

ing our RVM-fits geometrically, the inferred values of α from the normal/anomalous PA swings would correspond to the latitudinal positions of the two starspots on the neutron star, with the more active primary starspot positioned at $\alpha = 113 \pm 7^\circ$ and the secondary starspot located at $\alpha = 82 \pm 9^\circ$. Slight wobbles in the PA swings could be an indicator the distribution of magnetic field lines linking these two starspots is not uniform, while temporal variations in the polarisation profile could be due to a changing emission height and variable plasma flows along the coronal loop connecting the two starspots. Independent constraints on the viewing and emission geometry from X-ray observations, combined with continued radio monitoring would enable further tests of this hypothesis. Additionally, a simple comparison of the radio and X-ray profile alignment could test whether the radio emission originates from closed magnetic field lines above the hotspot or from open field lines at heights comparable to the light cylinder radius (see, e.g., [Camilo et al. 2007b](#) and [Gotthelf et al. 2019](#) for a discussion on the X-ray and radio profile alignment of XTE J1810–197).

6.7 Summary and conclusion

Our wide-band radio observations of Swift J1818.0–1607 have revealed the magnetar possesses highly active and dynamic magnetosphere following its 2020 outburst. This is highlighted by our detection of new profile components, and the appearance of transient emission and polarisation modes. We showed the post-outburst magnetic geometry remains stable across most of our observations, where variations in the linear PA and profile polarisation can potentially be ascribed to changes in the relative emission height over time. The reversed PA swing observed on MJD 59062 appears to be an anomalous outlier among our observations, which we speculate may be evidence of the radio emission at this epoch having originated from an additional, co-located magnetic pole that is offset from the primary pole by $\sim 30^\circ$ in latitude.

Continued monitoring of Swift J1818.0–1607 at radio wavelengths will allow for its magnetospheric evolution to be tracked as the current outburst progresses. This includes the detection of any new emission mode changing or deviations from the magnetic geometry that describes the majority of the data presented here. For instance, a series of high-cadence observations may be able to catch a transition from the normally positive sloping PA swing to the seemingly rare negative swing we observed on MJD 59062. Such a detection, combined with independent geometric constraints from fitting the X-ray profile and phase resolved spectrum of Swift J1818.0–1607, would provide an independent test

of the coronal loop hypothesis we proposed as a potential explanation for this phenomena.

7

Shining a light through a pulsar magnetosphere

The Double Pulsar PSR J0737–3039A/B is unique extra-solar laboratory. Consisting of two neutron stars bound together in a highly relativistic orbit that is remarkably edge-on from our perspective on Earth, this system has enabled stringent tests of more aspects of Einstein’s theory of General Relativity in the strong-gravity regime than any other Galactic double neutron star system to date. The high inclination of the binary orbit results in a brief ~ 30 second long eclipse of the fast-spinning pulsar A when it passes behind the slow-rotating pulsar B. Building upon the successful modelling of the eclipses by [Breton et al. \(2008\)](#), we characterised the secular evolution of the eclipse morphology due to the geodetic precession of pulsar B to obtain an updated precession rate of $\Omega_{\text{SO}}^{\text{B}} = 4.98^{+0.39}_{-0.35} \text{ yr}^{-1}$ (68% credible interval). This measurement is consistent with predictions from General Relativity to a relative uncertainty of 7.5%. Our analysis of the polarisation properties of pulsar A throughout the eclipse region revealed a significant amount of conversion between linear and circular polarisation is induced within the closed-field medium of B. The birefringent nature of the medium provides strong evidence the plasma contained within neutron star magnetospheres are indeed comprised of electron-positron pairs. We speculate the sign-changes in the handedness of the emergent circular polarisation likely result from changes in the projected, line-of-sight magnetic field direction of the B pulsar.

7.1 Introduction

PSR J0737–3039A/B is a highly relativistic double neutron star binary with a short 2.45 hr, mildly eccentric ($e = 0.088$) orbit ([Burgay et al., 2003](#)). Uniquely, both neutron stars have been detected as radio pulsars (referred to as pulsars A and B hereafter) with respective spin-periods of 22.7 ms and 2.8 s respectively ([Lyne et al., 2004](#)). High-precision timing of the two pulsars resulted in four independent tests of General Relativity (GR)

in the strong-field regime within only 2.7 yr since its initial discovery, making the system one of the most successful laboratories for testing our theories of gravity to date (Kramer et al., 2006b). The orbital plane of the system is inclined at an angle of $i = 89.35^\circ \pm 0.05^\circ$ (Kramer et al. submitted), remarkably edge-on from our perspective on Earth. This results in a 30-40 s long eclipse of pulsar A by the magnetosphere of pulsar B around A's superior conjunction (Lyne et al., 2004). The duration of these eclipses corresponds to a region of space that is $\sim 1.7 \times 10^7$ m wide, which spans only $\sim 10\%$ the light-cylinder radius of an equivalent isolated pulsar with the same rotational properties as pulsar B (Kaspi et al., 2004; Breton et al., 2012). This smaller than expected eclipsing region arises from the relativistic wind from pulsar A penetrating deep into the magnetosphere of pulsar B. As a result, the 'windward' side facing pulsar A is compressed, while the 'leeward' side is blown backwards into a cometary magnetotail (Arons et al., 2005).

Initial theoretical explanations for the observed phenomenology of the eclipse centred around the absorption of radiation from pulsar A by plasma embedded in the magnetosheath surrounding pulsar B's truncated magnetic field (Arons et al., 2005; Lyutikov, 2005). However, high-time resolution observations of the eclipses taken by McLaughlin et al. (2004) with the Robert C. Byrd Green Bank Telescope (GBT) in West Virginia, USA, revealed the light curve of pulsar A has peaks and troughs in its flux that are modulated at both once and twice the 2.8 s rotation period of pulsar B at different eclipse phases. This modulation in the light curve of pulsar A can be entirely explained through a simple geometric model, in which the radio pulses undergo synchrotron absorption by a relativistic pair-plasma that is confined to the toroidal, closed-field region of B's magnetosphere (Lyutikov & Thompson, 2005). The success of the model provided not only the first direct evidence for a dipole magnetic field geometry around a pulsar, but was later used to model a set of eclipse observations by the GBT over 3.9 yr, resulting in a novel detection of the geodetic precession rate of Pulsar B ($\Omega_{\text{GO}}^{\text{B}}$) and an associated fifth independent test of GR (Breton et al., 2008). While the effects of geodetic precession have been detected in five other relativistic binaries (Kramer, 1998; Kirsten et al., 2014; Fonseca et al., 2014; Venkatraman Krishnan et al., 2019; Desvignes et al., 2019), the corresponding precession rate measurements are largely indirect, primarily based on modelling the pulse profile width and polarisation properties. Interpretation of these observables is heavily reliant on the assumed pulsar beam-shape and the applicability of the rotating vector model (Radhakrishnan & Cooke, 1969), both of which are highly uncertain. Whereas Breton et al. (2008) showed the precession of pulsar B has a significant impact on the eclipse light curve over time, which was fitted by adding a simple linear drift in the spin-axis longitude

of pulsar B over time when computing the model templates.

In addition to enabling a unique test of GR, the eclipses provide a rare opportunity to directly probe the immediate plasma environment around a neutron star. High-sensitivity observations of the polarised radio pulses from pulsar A during and shortly after an eclipse can provide an insight into the largely unknown physical processes that take place within the closed-field region of B’s magnetosphere. Variations in the polarisation of A could also be used to infer the local magnetic field direction of B at different phases within the eclipse region, thereby providing an independent test of the Lyutikov & Thompson (2005) model.

7.2 MeerKAT observations

We have performed monthly monitoring observations of the Double Pulsar from July 2019 to May 2021 with the Meer Karoo Array Telescope (MeerKAT) in the Northern Cape province of South Africa under the MeerTime large science project. Observations at MeerKAT from March 2019 to March 2020 were performed using the 1284 MHz central frequency L-band receiver system, after which the majority of observations were performed with the 816 MHz central frequency UHF receivers. Data that are coherently dedispersed to account for the frequency-dependent delay induced by the passage of pulsar A’s radiation through the interstellar medium were collected using the PTUSE instrument (Bailes et al. 2020). PTUSE provides 1024 frequency channel filterbank data across the 856 MHz and 544 MHz bandwidths of the L-band and UHF receiver fleets respectively, along with $\sim 9 \mu\text{s}$ time sampling, and full Stokes information (Bailes et al., 2020). These data were then folded at the predicted rotation period of pulsar A using the DSPSR software package (van Straten & Bailes, 2011). Frequency channels that were significantly affected by radio-frequency interference were excised using the MEERGUARD¹ pulsar data cleaning package. We then binned the data in time by four rotations of pulsar A or a maximum time resolution of ~ 91 ms. The data were polarisation calibrated following the method outlined in section 3 of Serylak et al. (2021) and corrected for a rotation measure of $+120.82 \text{ rad m}^{-2}$ (Kramer et al., 2021).

7.3 Eclipse light-curve modelling

To model the total intensity light curve, we first extracted the flux of pulsar A via a matched-filtering process. This involved cross-correlating a high signal-to-noise template

¹<https://github.com/danielreardon/MeerGuard>

(generated from the integrated pulse profile from many hours of observations) with a frequency and polarisation averaged copy of the data. The resulting flux densities were normalised by the median off-eclipse value so the points where pulsar A is visible have values of order unity. Orbital phases for each flux point were derived by first converting the pulse arrival times at MeerKAT to equivalent arrival times at the Solar System barycenter, then computing the corresponding orbital phase using the pulsar A ephemeris of Kramer et al. (submitted). From here we then fitted the total-intensity data using a variation of the recipe outlined in Breton et al. (2008) in order to infer the geometry of pulsar B.

7.3.1 The Lyutikov and Thompson light-curve model

Modelling of the individual eclipse light curves was performed using the model outlined in Lyutikov & Thompson (2005). Here we briefly summarise the key features of the model and associated coordinate system.

A graphical overview of the model is shown in Figure 7.1. Pulsar B is positioned at the centre of a Cartesian coordinate system in which where the x -axis points toward our line-of-sight to the system, the z -axis is in the plane of the sky as viewed from Earth, and the y -axis runs parallel to the apparent motion of pulsar A (Breton et al., 2008). The motion of pulsar A is offset vertically from the origin by a constant value of $z = z_0$. The direction of Ω_B can be fully described by two angles: the spin-axis co-latitude (δ) with respect to the z -axis, and spin-axis longitude (φ_{so}) with respect to the x - y -plane (Damour & Taylor, 1992). Ordinarily, δ and φ_{so} would be related to our Cartesian coordinate system via

$$\begin{aligned}\theta &= \cos(90^\circ - i) \cos \delta - \sin(90^\circ - i) \sin \delta \cos \varphi_{so}, \\ \varphi &= \frac{\sin \delta \sin \varphi_{so}}{\sin \theta}.\end{aligned}\tag{7.1}$$

However, the Double Pulsar is nearly edge-on from our perspective ($i = 89.35 \pm 0.05^\circ$; Kramer et al. submitted), meaning the z -axis is effectively aligned with the total angular momentum vector. As a result, we can make the simplifying assumption that $\theta \approx \delta$ and $\varphi \approx \varphi_{so}$.

The radio intensity models of pulsar A throughout the eclipses are computed from the synchrotron optical depth (τ) of the plasma trapped within the closed-field region of pulsar B's toroidal magnetosphere. Following the prescription of Breton et al. (2008), the optical depth is at a given point during the eclipse is given by

$$\tau = \frac{\mu}{\nu_{\text{GHz}}^{-5/3}} \int_{-R_{\text{mag}}}^{+R_{\text{mag}}} \left(\frac{B \sin \kappa}{B_{\text{mag}}} \right) d\left(\frac{x}{R_{\text{mag}}} \right).\tag{7.2}$$

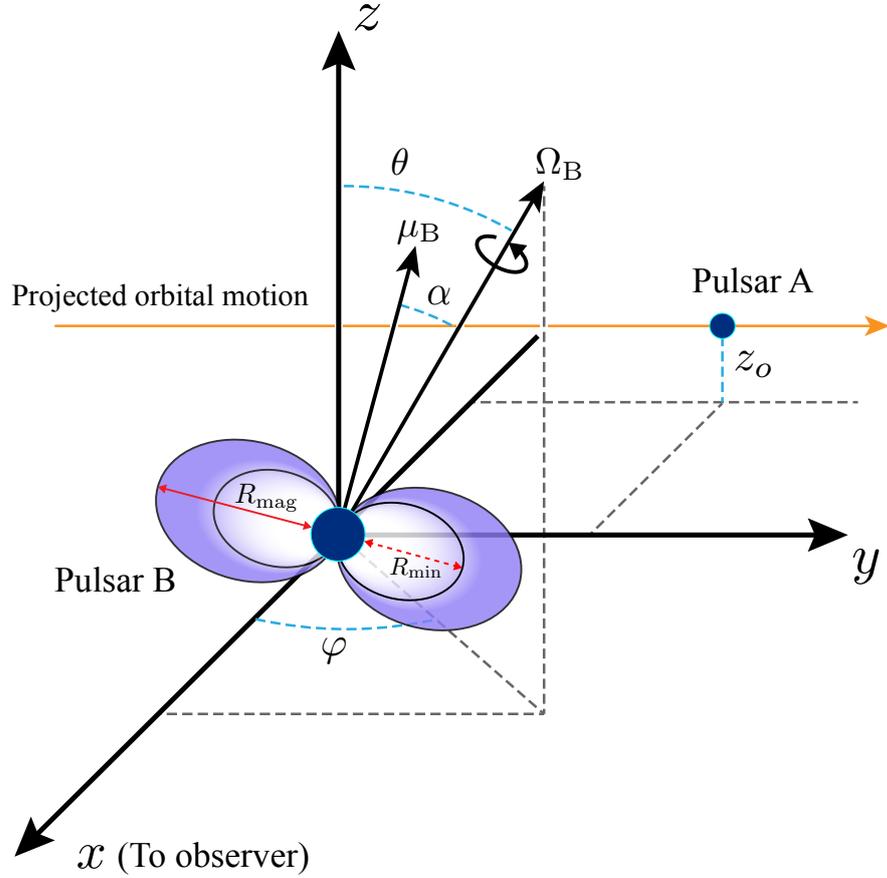


Figure 7.1 Diagram depicting the Double Pulsar geometry according to the eclipse light-curve model, adapted from figure 1 of [Breton et al. \(2008\)](#). Pulsar B is positioned at the origin of the Cartesian coordinate system. The projected orbital motion of pulsar A through superior conjunction is parallel to the y-axis. The closed, absorbing region of pulsar B's magnetosphere is represented by the shaded purple region truncated at R_{mag} . The minimum radial extent of the synchrotron absorbing plasma is represented by R_{min} . Longitude and latitude of the spin axis of pulsar B (Ω_B) is given by the angles φ and θ respectively. The magnetic axis (μ_B) is offset from Ω_B by angle α .

The terms in this equation are described as follows. The observing frequency (in GHz) is given by ν , R_{mag} is the truncation radius of B's magnetosphere, B is the local magnetic field strength along the line of sight in units of B_{mag} (magnetic-field strength at R_{mag}), κ is the angle between the local magnetic field strength and our line of sight and x is the radial position of pulsar A from our perspective in units of R_{mag} . The scaling parameter μ combines various parameters that describe the physical properties of B's magnetosphere as

$$\mu = \frac{4.5 \times 10^{-6} \lambda_{\text{mag}} k_B T_e}{N_B^{1/4} 10 m_e c^2}. \quad (7.3)$$

in which λ_{mag} is the pair plasma electron multiplicity, N_B alters the size of the magnetosphere based on the impact of the wind from pulsar A, k_B is the Boltzmann constant, T_e and m_e the electron temperature within the plasma, and c is the vacuum speed of light. Variations in each of these parameters, which are not practicable to be fit individually, serve to alter the depth of the eclipses as the intensity of pulsar A is computed from the optical depth as $e^{-\tau}$. Note that R_{mag} is not fit for directly but is inferred via the parameter ξ that scales the size of the magnetosphere to the orbital distance between pulsars A and B (Breton, 2009).

The modulation pattern of the light curve depends strongly on the changing line-of-sight geometry of pulsar B as it rotates, which is modelled through the corresponding variations in both B and κ in Equation 7.2. Both of these terms are related to the dipole unit vector magnetic polar angle (θ_μ) as

$$\cos \theta_\mu = \frac{\hat{\mu} \cdot \mathbf{r}}{r}, \quad (7.4)$$

where $\mathbf{r} = \{x, y(t), z\}$, r is the distance between pulsars A and B in spherical coordinates ($r = |\mathbf{r}| = \sqrt{x^2 + y^2(t) + z^2}$) and $\hat{\mu}$ is the dipole unit vector, the components of which are given by

$$\begin{aligned} \hat{\mu}_x &= (\hat{\mu}_x^\Omega \cos \theta + \hat{\mu}_z^\Omega \sin \theta) \cos \varphi - \hat{\mu}_y^\Omega \sin \varphi, \\ \hat{\mu}_y &= (\hat{\mu}_x^\Omega \cos \theta + \hat{\mu}_z^\Omega \sin \theta) \sin \varphi + \hat{\mu}_y^\Omega \cos \varphi, \\ \hat{\mu}_z &= \hat{\mu}_z^\Omega \cos \theta - \hat{\mu}_x^\Omega \sin \theta, \end{aligned} \quad (7.5)$$

with

$$\begin{aligned} \hat{\mu}_x^\Omega &= \sin \alpha \cos(\phi_B + \Delta\phi_B), \\ \hat{\mu}_y^\Omega &= \sin \alpha \sin(\phi_B + \Delta\phi_B), \\ \hat{\mu}_z^\Omega &= \cos \alpha. \end{aligned} \quad (7.6)$$

Here, α is the magnetic inclination angle of pulsar B and ϕ_B is the rotation phase of pulsar B, which is related to the spin-vector direction as $\phi_B = \Omega_B t = 2\pi t/P_B$ where P_B is the spin period of the pulsar. The parameter $\Delta\phi_B$ accounts for the offset from $\phi_B = 0$ at an assumed reference time. The values of B and κ are computed at each step in \mathbf{r} as

$$B = \frac{\sqrt{1 + 3 \cos^2 \theta_\mu}}{r^3} \mu_B, \quad (7.7)$$

and

$$\cos \kappa = \frac{3 \cos \theta_\mu (x/r) - \hat{\mu}_x}{1 + 3 \cos^2 \theta_\mu}, \quad (7.8)$$

before being passed to Equation 7.2 where they are integrated over in terms of x .

Deriving pulsar B arrival times from the eclipses

The time shift between the profile and the template yields the TOA relative to a fiducial point of the template and the start time of the observation (Taylor 1992)

In pulsar timing, the arrival time of a pulse is defined by the relative offset between a fiducial reference point in the pulse profile and the start of the observation (Taylor, 1992). For a single-component pulse profile this point typically corresponds to the profile peak. Since the radio pulses from pulsar B are currently undetected, we instead define the fiducial reference point to be when the magnetic moment of the pulsar is maximally pointed in the direction of the Earth. In the Lyutkiov and Thompson model, this occurs when $\hat{\mu}_x$ is at a maximum, i.e

$$\max(\hat{\mu}_x) = \frac{d\hat{\mu}_x}{d\phi_B} = 0. \quad (7.9)$$

Substituting in the appropriate terms from Equations 7.6 and 7.5, the additional phase-shift required to maximise $\hat{\mu}_x$ ($\delta\phi_B$) can be found as

$$\begin{aligned} \sin \varphi \cos(\delta\phi_B) + \cos \theta \sin(\delta\phi_B) \cos(\varphi) &= 0 \\ \Rightarrow \delta\phi_B &= \tan^{-1} \left(\frac{-\sin \varphi}{\cos \varphi \cos \theta} \right), \end{aligned} \quad (7.10)$$

with the maximum in $\hat{\mu}_x$ occurring when

$$\frac{d^2\hat{\mu}_x}{d\phi_B^2} = \sin(\delta\phi_B)(\sin \varphi \sin(\delta\phi_B) - \cos \theta \cos(\delta\phi_B) \cos \varphi) < 0. \quad (7.11)$$

This is implemented in the model by adding the result of Equation 7.10 to obtain the $\Delta\phi_B$ phase-offset to determine the total rotation required at each reference epoch such that μ_B is maximally pointed at the Earth. The combination of this offset, in addition to the rotation period of pulsar B, can in principle be combined to form an ‘effective’ ToA, as

$$\Delta t_{\text{SSB}} = \Delta t_{\text{ref}} + P_B(\Delta\phi_B + \delta\phi_B) \quad (7.12)$$

where Δt_{ref} is the Barycentric reference time of the eclipse data.

7.3.2 Joint-fitting of eclipse pairs

Unlike in [Breton et al. \(2008\)](#), radio pulses are not currently detected from pulsar B. Hence, we were unable to make use of a phase-connected timing solution to infer the precise rotation phase of the pulsar during our observations and could therefore perform a simultaneous joint-fit to every light curve within our sample. However, the S/N of the eclipses detected by MeerKAT is sufficiently high that we can instead perform direct fitting of individual eclipses, where the rotation-phase of pulsar B is included as a free parameter.

Our initial attempts at fitting the eclipses independently of one another were severely affected by epoch-to-epoch stochasticity in the eclipse envelope. This phenomena added a significant amount of scatter in the geometric constraints inferred from one eclipse to the next. It probably originates from random fluctuations of the plasma content and radial extent of the closed magnetic field lines of pulsar B. We were able to mitigate a significant amount of this behaviour by performing joint fits to pairs of eclipses that were separated in time by only a single orbit. This was conducted using a Gaussian likelihood function of the form

$$\mathcal{L}(\mathbf{d}|\Theta) = \prod_{i=1}^2 \prod_{j=1}^N \frac{1}{\sqrt{2\pi\hat{\sigma}_i^2}} \exp\left[-\frac{(d_{i,j} - \mu_{i,j}(\Theta))^2}{2\hat{\sigma}_i^2}\right], \quad (7.13)$$

where \mathbf{d} represents the input eclipse light curves, μ is the eclipse model, Θ contains the model parameters, and $\hat{\sigma}_i^2 = \sigma_i^2 + \sigma_{Q,i}^2$ are the uncertainties on the light-curve fluxes added in quadrature with an additional error parameter ($\sigma_{Q,i}$). This extra uncertainty parameter accounts for both pulse-to-pulse flux variations of pulsar A and unaccounted systematic errors. Any drift in the expected rotation-phase of pulsar B between these pairs of eclipses would be minimal, and we therefore specified only a single $\Delta\phi_B$ parameter to infer the phase-shift in pulsar B's rotation at the reference time associated with the first eclipse. Our priors for the model parameters are summarised in [Table 7.1](#). The spin period of pulsar B was fixed to the predicted value at each observing epoch. Posterior samples for the model parameters were generated using BILBY as a front-end to the PYMULTINEST sampler, a PYTHON-based implementation of the MULTINEST nested-sampling algorithm ([Ashton et al., 2019b](#); [Buchner et al., 2014](#); [Feroz et al., 2009](#)).

The joint-modelling of eclipse pairs were able to reproduce most of the observed eclipses phenomenology detected by MeerKAT at both L-band and UHF frequencies (see [Figure 7.2](#)). Similar to [Breton et al. \(2008\)](#), the eclipse fits are poorest near the ingress and egress phases, where the effects of small variations in magnetosphere size or plasma density have the largest impact on the observed light curve. However, we found that excluding the

Table 7.1 Priors on the eclipse light curve parameters. Parameters with Gaussian priors are emulate the reported values and uncertainties of [Breton et al. \(2008\)](#).

Parameter	Symbol	Prior type
Pulsar B phase offset	$\Delta\phi$	Uniform(0, 1)
Magnetic inclination angle ($^\circ$)	α	Gaussian(130, 1)
Spin axis co-latitude ($^\circ$)	θ	Gaussian(71, 1)
Spin axis co-longitude ($^\circ$)	φ	Uniform(-90, 90)
Minimum synchrotron absorption radius	r_{\min}	Uniform(0, 1)
Characteristic optical depth	μ	Gaussian(2, 0.2)
Eclipse scale factor	ξ	Gaussian(1.29, 0.2)
Vertical offset of pulsar A	z_0/R_{mag}	Gaussian(-0.543, 0.1)

egress by restricting the orbital extent of the eclipse light curves to phases ranging from -1.0° - 0.75° had no impact on the recovered light-curve model. This may be related to the modulation pattern in the light curve having become more symmetric due to the spin-precession of pulsar B.

Despite the relative success of the joint-fits, we found systematic offsets existed between the geometric parameters measured using the L-band and UHF receivers. These offsets persisted regardless of whether we applied additional offsets to the orbital phases to account for excess dispersion delays or relativistic aberration prior to re-running the inference framework on both synthetic and real datasets. They are also unlikely to be the result of ignoring higher-order multipole components within the magnetosphere of pulsar B, as the rapid increase in optical depth means the eclipses are restricted to probing the far-field regime. Additionally, [Breton et al. \(2008\)](#) noted that attempts to fit the eclipses with a purely quadrupolar magnetic field structure failed to accurately recover the observed modulation pattern. Hence it is likely to originate from an incorrect frequency-dependence within the light-curve model. Comparing the median values between the two bands, we inferred offsets of $\delta\alpha = 3.6^\circ$, $\delta\theta = 12.6^\circ$ and $\delta\varphi = 4.8^\circ$. Figure 7.3 shows the recovered posteriors of these three parameters over time after removing their respective offsets. The systematic offsets also impacted our extraction of effective ToAs for pulsar B, and we were unable to obtain a phase-connected timing solution for pulsar B across our 1.8 yr observing span. Part of this resulted from the joint-fits effectively halving the total number of possible ToAs that could be extracted. Additionally, our approximately monthly observing cadence makes it difficult to account for possible phase-wraps in the timing residuals. Both of the systematic issues with the modelling and cadence obstacle can be overcome

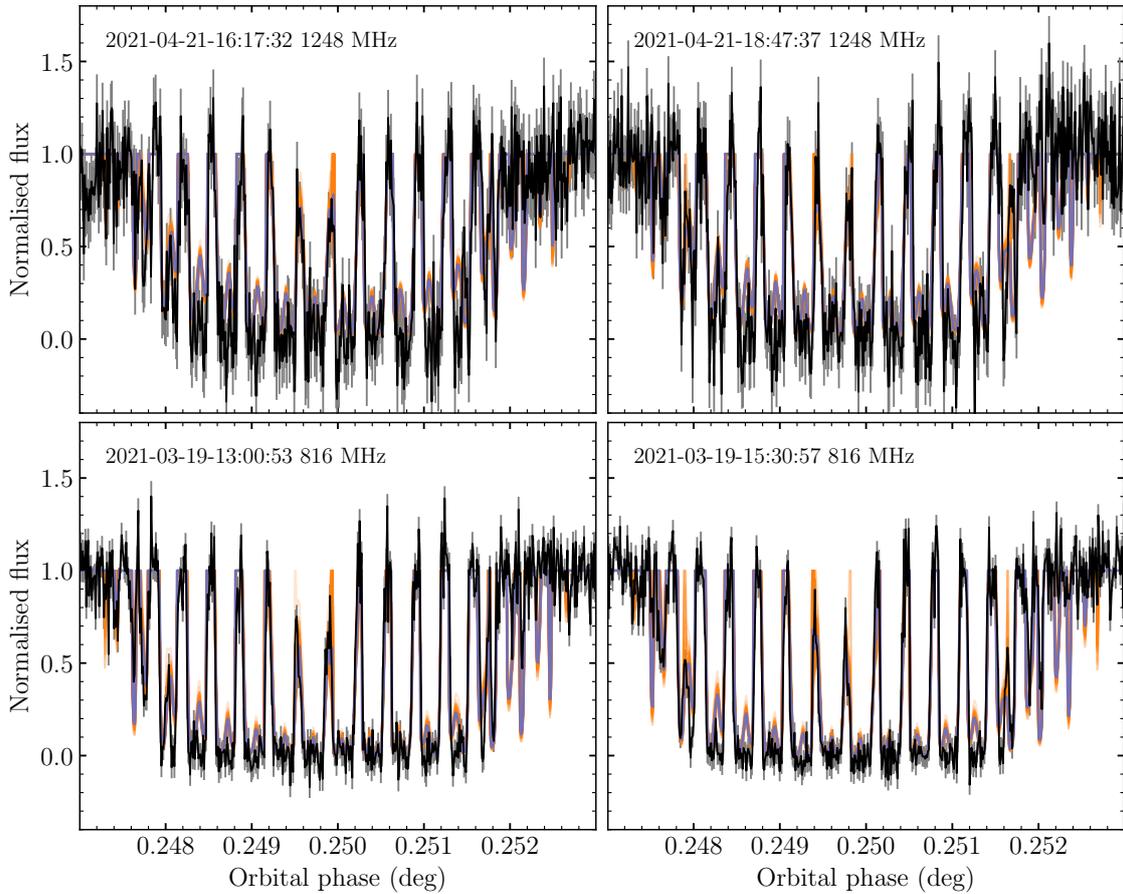


Figure 7.2 Two pairs of pulsar A eclipse profiles detected during two separate observing epochs (black) with the median a-posteriori recovered light curve (purple) and 100 random draws from the posterior distributions (orange) over-plotted.

through improvements to the frequency dependence of the [Lyutikov & Thompson \(2005\)](#) model, potentially through expanding it to fit both time and frequency information simultaneously, and by performing a short observing campaign to obtain a dense set of eclipses within a small timespan. The latter would guarantee a phase-connected timing solution could be obtained once the outstanding systematic effects are fully accounted for.

7.3.3 Spin-orbit precession of pulsar B

The spin-vector of a rotating body moving in the curved spacetime of a companion will precess about the total angular-momentum vector of the binary system. In the case of the Double Pulsar, the total angular momentum vector is essentially aligned with the orbit-normal vector (z -axis in [Figure 7.1](#)). As a result, the geodetic precession of pulsar B

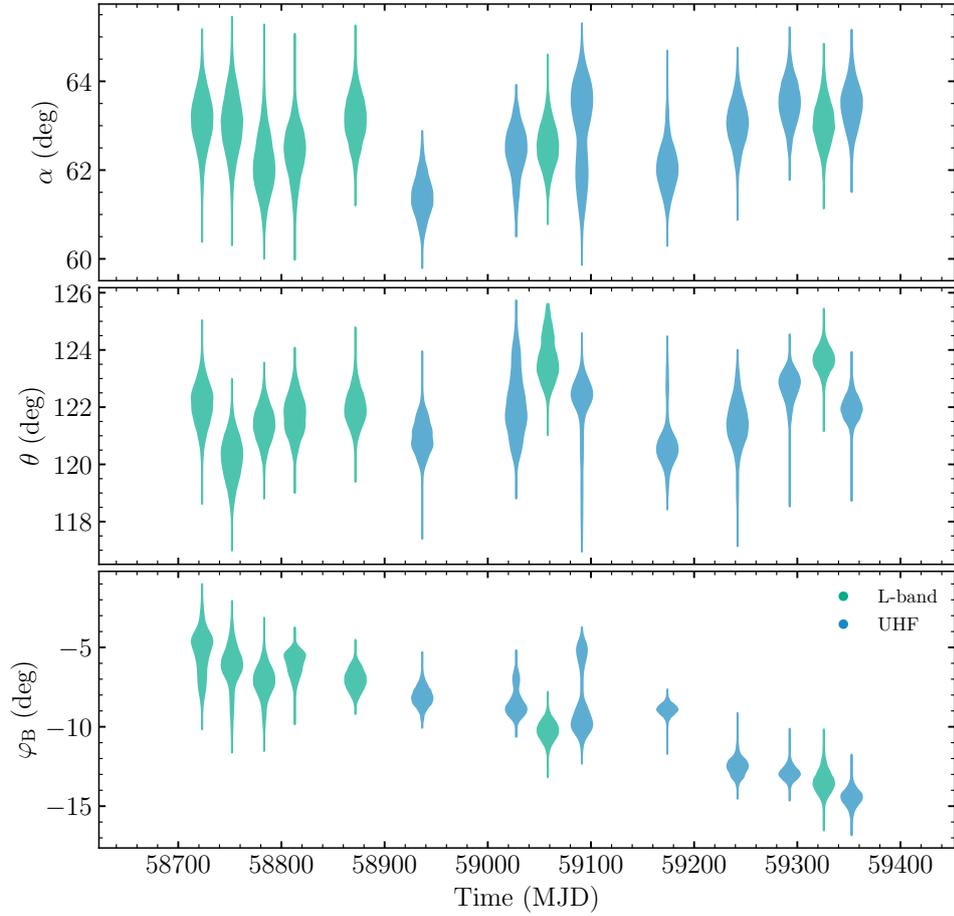


Figure 7.3 Evolution of pulsar B’s geometry over time. Posterior distributions of α (top), θ (middle) and φ after subtracting the systematic offsets between observing bands. Green shading corresponds to measurements from L-band observations and blue from UHF.

manifests as a time-varying change in the spin-axis longitude while the spin-axis latitude remains unchanged. Hence, the time-evolution of these two parameters can be written as

$$\theta = \theta_0, \quad (7.14)$$

and

$$\varphi(t) = \varphi_0 - \Omega_{\text{SO}}^{\text{B}}(t - t_0), \quad (7.15)$$

where θ_0 and φ_0 are the co-latitude and longitude of B’s spin axis at some reference epoch, t_0 . This is highlighted in Figure 7.3, where θ shows no evidence of any time-evolution over the 1.8 yr observing span, while φ shows a clear linear drift with a negative gradient.

As a phase-connected timing solution for pulsar B across our dataset is unavailable,

we could not use the same simultaneous joint-fit method utilised by [Breton et al. \(2008\)](#) to infer the precession rate of pulsar B. Instead, we performed a second-stage hierarchical fit to the posterior samples we obtained for φ at each observing epoch. We used a hyper-likelihood function of the form

$$\mathcal{L}_{\text{tot}}(\varphi_0, \Omega_B | \varphi(t)) = \prod_{i=1}^{N_e} \frac{1}{n_i} \sum_k^{n_i} \exp \left[- \frac{(\varphi(t_i)_k - \varphi_0 + \Omega_B (t_i - t_0))^2}{2\sigma_\varphi^2} \right], \quad (7.16)$$

where N_e is the number of epochs for which we have measured φ , n_i is the total number of posterior samples from the i -th eclipse, and σ_φ is a normalising factor that accounts for the variance in our measurements of φ . Uniform priors were assumed for each the hyper-parameters. As with the joint-eclipse fits, the hyper-posteriors were sampled using the PYMULTINEST sampler with BILBY. Fits to the L-band and UHF results were performed independent of one another to avoid introducing any systematic bias from an imperfect subtraction of the offset in φ between the two bands.

The marginalised posterior distributions for our L-band and UHF measurements of $\Omega_{\text{SO}}^{\text{B}}$ are compared in [Figure 7.4](#). Individually, we recovered a precession rate of $\Omega_{\text{SO}}^{\text{B}} = 5.02^\circ_{-0.50^\circ}^{+0.52^\circ} \text{ yr}^{-1}$ from the L-band posterior samples and $\Omega_{\text{SO}}^{\text{B}} = 5.00^\circ_{-0.61^\circ}^{+0.60^\circ} \text{ yr}^{-1}$ with the UHF results. Taking the product of the posterior distributions inferred from both bands, we obtained a combined measurement of $\Omega_{\text{SO}}^{\text{B}} = 4.98^\circ_{-0.35^\circ}^{+0.39^\circ} \text{ yr}^{-1}$. This measurement agrees with the expected value of $5.0734^\circ \pm 0.0007^\circ \text{ yr}^{-1}$ from GR to within ~ 7.4 percent, and represents almost a factor of two improvement over the results of [Breton et al. \(2008\)](#).

The improved measurement of the geodetic precession rate can also be used in combination with measurements of other relativistic effects to constrain alternate theories of gravity. In [Figure 7.5](#), we have plotted the inferred masses of pulsars A and B from $\Omega_{\text{SO}}^{\text{B}}$ alongside those from post-Keplerian parameters measured through precision timing of pulsar A ([Kramer et al. submitted](#)) and the mass-ratio ([Kramer et al., 2006a](#)). The pair of lines associated with the 68% credible interval for $\Omega_{\text{SO}}^{\text{B}}$ intersects the same common-point as the other post-Keplerian parameters, indicating that it is consistent with GR to our current measurement uncertainty. In addition to the mass-mass comparison, [Breton et al. \(2008\)](#) also demonstrated that measurements of $\Omega_{\text{SO}}^{\text{B}}$ combined with the timing of both pulsars provides the only relativistic binary system in which direct constraints can be placed on the strong-field spin-orbit precession. Under a set of generic Lorentz-invariant relativistic theories introduced by [Damour & Taylor \(1992\)](#), the geodetic precession rate can be reformulated as

$$\Omega_{\text{SO}}^{\text{B}} = \left(\frac{P_b}{2\pi} \right)^{-3} \frac{x_A x_B}{s^2(1-e^2)} \frac{c^2 \sigma_{\text{so}}}{G}, \quad (7.17)$$

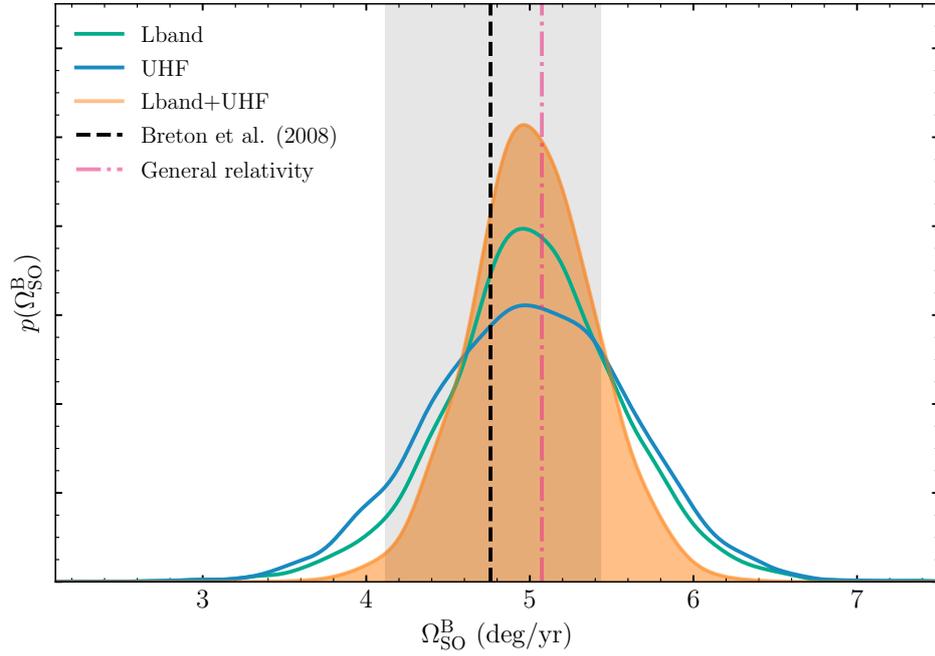


Figure 7.4 Posterior distributions of $\Omega_{\text{SO}}^{\text{B}}$ inferred from independent fits to the L-band (green) and UHF (blue) data, and the combined posterior from both (orange). Median (dashed line) and 68% credible intervals from [Breton et al. \(2008\)](#) are shown black, while the predicted value from General Relativity is given by the dash-dotted magenta line.

where x_{A} and x_{B} are the projected semi-major axes as measured from timing of both pulsars A and B, $s = \sin i$ is the Shapiro-delay shape parameter, e is the orbital eccentricity, P_b the orbital period, c the vacuum speed of light, σ_{so} is the spin-orbit coupling constant and G is a generalised gravitational constant for the interaction between the two pulsars. If GR is the correct theory of gravity, then we expect $\left(\frac{c^2\sigma_{\text{so}}}{G}\right)_{\text{GR}} = 2 + \frac{3}{2}\frac{m_{\text{A}}}{m_{\text{B}}} = 3.607274 \pm 0.000022$, using the masses of pulsars A and B from the intersection of the relativistic periastron advance ($\dot{\omega}$) and Shapiro-delay shape ([Kramer et al. submitted](#)). Substituting in our value of $\Omega_{\text{SO}}^{\text{B}}$ and measurements for the post-Keplerian parameters ([Kramer et al. 2006a](#), submitted) into [Figure 7.17](#), we obtain $\left(\frac{c^2\sigma_{\text{so}}}{G}\right) = 3.54 \pm 0.27$. Taking the ratio of the observed and predicted values, we find $\left(\frac{c^2\sigma_{\text{so}}}{G}\right)_{\text{obs}} / \left(\frac{c^2\sigma_{\text{so}}}{G}\right)_{\text{GR}} = (1 - 0.019) \pm 0.075$. Hence, our measured spin-orbit coupling constant is consistent with the expectation from GR to within an uncertainty of 7.5 percent. The level of consistency between this test and the comparison between the predicted and observed precession rates from earlier is expected as both are reliant on the same measured value.

While these measurements of the geodetic precession rate of pulsar B and the associated test of spin-orbit coupling are encouraging, we have not attempted to quantify the impact

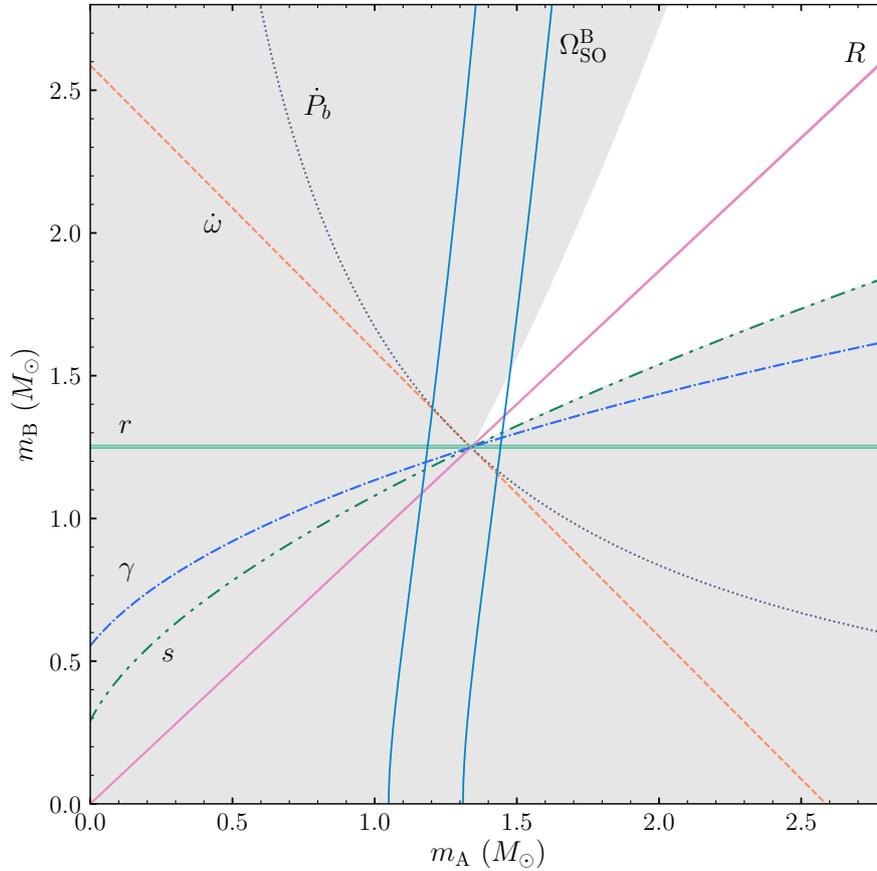


Figure 7.5 Mass-mass diagram illustrating current tests of General Relativity with the Double Pulsar. Shaded region is forbidden by the individual mass functions of the two pulsars, as $\sin i \leq 1$. Constraints on the masses of pulsar A and B are shown by pairs of lines, where the separation between them indicates the $1 - \sigma$ uncertainty. If General Relativity is the correct theory of gravity, then all lines should intersect at a common point.

of additional systematic effects that arise from the aforementioned inconsistencies between the model and the data. [Breton et al. \(2008\)](#) noted that variations in the input rotational phase of pulsar B can lead to slightly faster or slower precession rates being inferred. As a result, the values presented here may have somewhat underestimated uncertainties. Ultimately, this will be addressed through future alterations to the [Lyutikov & Thompson \(2005\)](#) model that will result in it better matching the observed eclipse phenomenology at the two different observing bands.

7.4 Spectropolarimetry of the eclipses

In addition to enabling measurements of the geometric orientation of pulsar B’s magnetosphere, the eclipses also provide a novel means to directly probe the plasma trapped within the closed-field region through studying the polarimetry of pulsar A. Faraday rotation occurs when linearly polarised radiation propagates through a medium with differential, circularly polarised natural wave-modes. The difference in refractive index between the two wave-modes results in a phase-delay between the right- and left-handed polarisations of the transmitted radiation that we detect as a rotation of the linear polarisation vector. However, in the case of a more general birefringent medium, the natural wave-modes can be either elliptically or linearly polarised. This results in a process known as ‘Faraday conversion’, a generalised form of Faraday rotation where a conversion between linearly and circularly polarised radiation takes place. The success of the [Lyutikov & Thompson \(2005\)](#) eclipse model suggests the magnetospheric plasma surrounding pulsar B is comprised of a highly relativistic, electron-positron pair-plasma. One may expect the electrons and positrons to contribute Faraday conversion with equal and opposite signs, thereby cancelling it out. However, slight asymmetries in the ratio of one particle over the other can result in a significant amount of detectable circular polarisation being produced ([Sazonov, 1969](#); [Noerdlinger, 1978](#)). As a result, we may expect to detect some amount of Faraday conversion in the polarised radiation from pulsar A that propagates the edges of the magnetosphere. Here the optical depth is lower, resulting in only a partial absorption of the incident radiation. Additional Faraday rotation may also occur immediately following the eclipse region if the radio pulses from pulsar A interact with cooler, mildly-relativistic particles within the magnetotail of pulsar B.

7.4.1 Polarised light curves

To search for variations in the polarisation properties of pulsar A, we first extracted its polarisation light curves from a set of 21 polarisation calibrated MeerKAT observations (13 at L-band, 8 at UHF). We extracted the light curves by averaging the flux from pulsar A within two windows centred on its two profile components, as shown in [Figure 7.6](#). Associated orbital phases for each point were computed in the same fashion as [Section 7.3](#). The individual light curves were converted to Stokes format, where Stokes I is the total intensity, Q and U the two linear polarisations (total linear polarisation is given by $L = \sqrt{Q^2 + U^2}$), and V is the circular polarisation. The handedness of Stokes V follows the pulsar/IEEE convention, where positive values correspond to right-

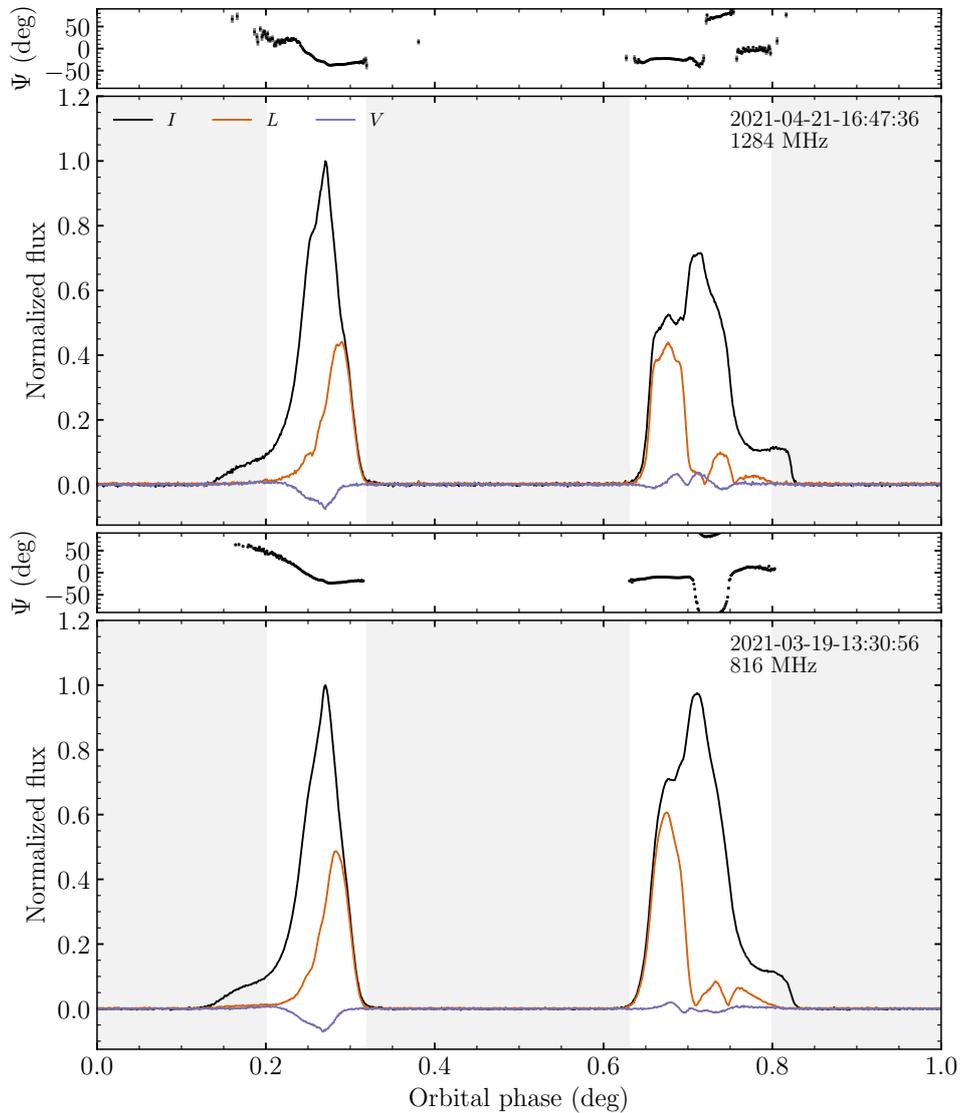


Figure 7.6 Example time and frequency-averaged pulsar A polarisation profiles detected by MeerKAT at L-band (top) and UHF (bottom). The upper panels show the linear polarisation position angle (Ψ), while the lower panels depict the total intensity (black), linear polarisation (orange) and circular polarisation (purple) profiles. The grey-regions indicate pulse phases that were excluded when extracting the polarisation light curves.

hand circular (RHC) polarisation and negative values are left-handed (LHC). As with the matched-filtered flux measurements, we normalised the polarisation light curves by the median off-eclipse total-intensity value. To increase our sensitivity to propagation effects introduced by the ‘bulk’ of pulsar B’s magnetosphere, we averaged the polarisation light curves in each band and then binned them in orbital phase. Each of the resulting orbital

phase bins covers $\sim 0.0625^\circ$ or ~ 1.5 s. The polarisation light curves, total polarisation ($\Pi = \sqrt{Q^2 + U^2 + V^2}/I$), polarisation position angle (PA; $\Psi = 0.5 \tan^{-1}(U/Q)$) measurements and RM fits are displayed in Figure 7.7. The uncertainties in the PA at each orbital phase were computed as $\sigma_\Psi = 28.65^\circ(\sigma_I/L_c)$ (Lorimer & Kramer, 2004), where σ_I is the standard deviation of Stokes I and L_c is the noise de-biased linear polarisation (Everett & Weisberg, 2001). We also searched for deviations in the rotation measure (ΔRM) of pulsar A at each orbital-phase point. This was performed by first binning the data in frequency to 64 channels across the L-band and UHF bandpasses, followed by directly fitting the resulting Stokes Q and U spectra using the method outlined by Bannister et al. (2019).

7.4.2 Position angle variations and Faraday rotation

Both the UHF and L-band total intensity and linear polarisation light curves undergo similar decreases in intensity throughout the eclipse region, reaching a minimum just prior to superior conjunction. The average eclipse region appears slightly wider at UHF than at L-band, which is consistent with previous studies (e.g. Kaspi et al., 2004; Breton et al., 2012). A small peak in the L-band Stokes L curve near orbital phase 89.7° appears to have been associated with a $\Delta\Psi = 20^\circ \pm 4^\circ$ increase in the observed PA. This peak is followed by a possible dip near 89.8° , however the large PA uncertainties within this part of the eclipse region means it is difficult to distinguish it from a random fluctuation due to low S/N . Another small peak in the L-band PA also occurs during the egress phase with $\Delta\Psi = 13^\circ \pm 5^\circ$. There are at least three relative increases visible the UHF PAs at orbital phases close to 89.3° , 90.3° and 90.6° . The first peak has a relative change in PA of $\Delta\Psi = 7^\circ \pm 3^\circ$, while the two peaks during the egress phase have $\Delta\Psi =$ and $10^\circ \pm 3^\circ$ respectively. Similar PA variations were previously reported by Yuen et al. (2012) in observations taken by the Parkes telescope at 20- and 50-cm wavelengths, the amplitudes of which were suggested to be consistent with differential absorption of the linearly polarised components of pulsar A's radio pulses. From Lyutikov & Thompson (2005), the absorption coefficients for polarised radiation with components that are parallel and perpendicular to the line-of-sight magnetic field are $a_\nu^\parallel = 4a_\nu/3$ and $a_\nu^\perp = 2a_\nu/3$ respectively. Hence, there is a preferential absorption of radiation that has a polarisation direction that is perpendicular to the magnetic field direction. As a result, the emergent radio waves that have been transmitted through regions of increasing optical depth will result in transmitted PAs that approach the projected angle of pulsar B's magnetic field on the sky.

Changes in the PA following the eclipse region may arise from interactions between the radio pulses from pulsar A and mildly relativistic particles within the extended magnetotail

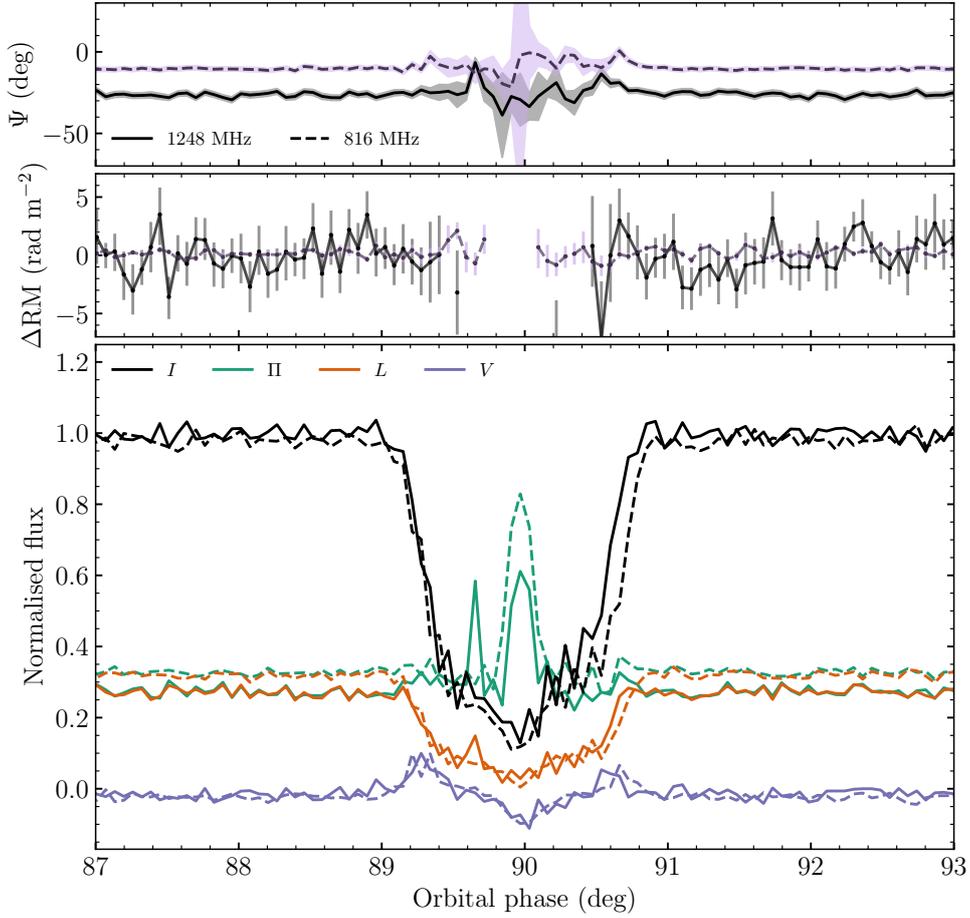


Figure 7.7 Orbital-phase averaged polarisation light curve of pulsar A at L-band (solid lines) and UHF (dashed lines). The top and middle panels show the linear polarisation position angle (Ψ) and deviation in the rotation measure of pulsar A (ΔRM) from $+120.84 \text{ rad}^{-2}$, where the uncertainties for the L-band data are shown in grey and UHF in purple. The bottom panel depicts the polarisation light curve with total intensity in black, total polarisation (Π) in green, linear polarisation in orange and circular polarisation in purple.

of pulsar B. [Yuen et al. \(2012\)](#) speculated the low-significance increases in PA within the egress phase of the Parkes telescope’s eclipses may have resulted from this effect. While the standard deviation of the plotted median a posteriori RM values in the post-eclipse regime is marginally higher than the pre-eclipse value ($\sigma_{\text{RM,med}} = 1.5$ versus $\sigma_{\text{RM,med}} = 1.8$ at L-band; $\sigma_{\text{RM,med}} = 0.32$ versus $\sigma_{\text{RM,med}} = 0.29$ at UHF), it is not significant enough to claim a detection of excess Faraday rotation. We also did not recover significant deviations in the RM of pulsar A at the same orbital phases as our detected increases in PA during the egress phase at both L-band and UHF frequencies.

7.4.3 Birefringence induced circular polarisation

In addition to the PA variations, we also observed increased levels of circular polarisation throughout the eclipse region with three discrete peaks occurring at different phases. An initial rise and peak is clearly visible in both bands during the early-ingress phase, which is followed by a sign-change indicating a flip from RHC to LHC that subsequently peaks at superior conjunction. The handedness of the polarisation then reverts back to RHC prior to a third and final peak during egress. We also detected a large spike in the total polarisation fraction in both bands. Such an increase in total polarisation was predicted to occur by [Lyutikov & Thompson \(2005\)](#), as unpolarised radiation is partially converted into polarised radiation as it propagates through the magnetosphere of pulsar B. The offset in the total polarisation peak when compared to figure 11 of [Lyutikov & Thompson \(2005\)](#) can be ascribed to the deepest part of the eclipse slowly evolving over time as the pulsar precesses (see figure 12 from their work). From Figure 7.7, it is clear the total polarisation peak coincides with the aforementioned increase in LHC polarisation as opposed to the predicted increase in linear polarisation. This difference likely arises from incident radiation from pulsar A being substantially linearly polarised (see Figure 7.6) as opposed to the assumed unpolarised background source in [Lyutikov & Thompson \(2005\)](#). The near constant total polarisation throughout the eclipse region (peak at superior conjunction notwithstanding) suggests the increased levels of circular polarisation are generated by a conversion of the incident linearly polarised radiation from pulsar A into circular polarisation via generalised Faraday rotation. This can occur if the propagating medium within pulsar B's magnetosphere is comprised of a highly birefringent pair-plasma.

Modelling the polarisation spectrum at these particular orbital phases can potentially reveal the underlying process that is responsible for the generalised Faraday rotation. Different birefringent media can produce polarisation spectra with differing spectral exponents that relate to the properties of the intervening particles, magnetic field strength and projected angle with respect to the line of sight (e.g. [Kennett & Melrose, 1998](#)). We fit the observed spectra normalised by the total polarisation (P) at each of the Stokes V peaks using a simple phenomenological generalised Faraday rotation model of the form

$$\hat{\mathbf{P}}(\lambda) = \begin{bmatrix} \cos(\vartheta) & 0 & \sin(\vartheta) \\ 0 & 1 & 0 \\ -\sin(\vartheta) & 0 & \cos(\vartheta) \end{bmatrix} \cdot \begin{bmatrix} \cos(\varrho) & -\sin(\varrho) & 0 \\ \sin(\varrho) & \cos(\varrho) & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos[2\Psi(\lambda)] \cos(2\chi_0) \\ \sin[2\Psi(\lambda)] \cos(2\chi_0) \\ \sin(2\chi_0) \end{bmatrix}, \quad (7.18)$$

where λ is the observing wavelength, ϑ and ϱ are the angles in which the polarisation vector is rotated in the Stokes U and V directions respectively, χ_0 is the ellipticity angle

$(0.5 \tan^{-1}(V/L))$ at our reference frequency, and $\Psi(\lambda)$ is the position angle given by

$$\Psi(\lambda) = \Psi_0 + \text{GRM}(\lambda^\varsigma - \lambda_c^\varsigma). \quad (7.19)$$

Here, Ψ_0 is the position angle at the reference frequency, GRM is the ‘generalised rotation measure’ and ς is the spectral exponent of the generalised Faraday rotation. The direction in which natural wave-modes of the birefringent medium are polarised can be inferred from ϑ , where ‘standard’ Faraday rotation arising from circularly polarised modes would result in $\vartheta = 0$ degrees and a spectral exponent of $\varsigma = 2$. Tilting of the polarisation plane away from the V-axis ($\vartheta > 0$) would indicate the presence of an elliptically polarised medium, with $\vartheta = 90$ degrees would result from a medium with purely linearly polarised natural wave modes (e.g. a pair-plasma).

Using PYMULTINEST, we sampled the posterior distributions of the model parameters with a Gaussian likelihood function of the form

$$\mathcal{L}(\mathbf{P}(\lambda)|\Theta) = \prod_i^N \frac{1}{\sqrt{2\pi}\sigma^2} \exp \left[- \frac{(\mathbf{P}(\lambda_i) - \hat{\mathbf{P}}_m(\Theta; \lambda_i))^2}{2\sigma^2} \right], \quad (7.20)$$

where $\mathbf{P}(\lambda)$ is the input polarisation spectra, $\hat{\mathbf{P}}_m(\lambda; \theta)$ is our generalised Faraday rotation model, and σ is a free parameter that approximates the variance of the data. Despite the large fractional bandwidths afforded by the MeerKAT receivers, our attempts to constrain the spectral dependence of the polarisation spectra returned values of $\chi_0 \sim \pm 45^\circ$ and $\varsigma \sim 0$, indicating the spectra are almost entirely circularly polarised and show no frequency-dependent conversion between the Stokes parameters. This can be seen in the normalised Stokes spectra plotted for each of the circular polarisation peaks in Figures 7.8, 7.9 and 7.10. There are several possible explanations for the lack of frequency dependence in the polarisation spectra at the circular polarisation peaks. It is possible that our averaging of several eclipses in orbital phase could have smoothed over small phase offsets in the PA caused by stochasticity in the pair-plasma density from one eclipse to the next. The rotation of pulsar B also induces a changing magnetic field strength and projected direction along the line of sight, that is somewhat averaged over due to our finite orbital-phase bin widths. Additionally, magnetic field reversals along the line of sight can induce a rapidly oscillating form of generalised Faraday rotation with a λ^2 scaling (e.g. Melrose, 2010; Gruzinov & Levin, 2019). This behaviour would also have been smeared out by our averaging over multiple eclipses.

A possible link to the magnetic field direction playing an important role in generating

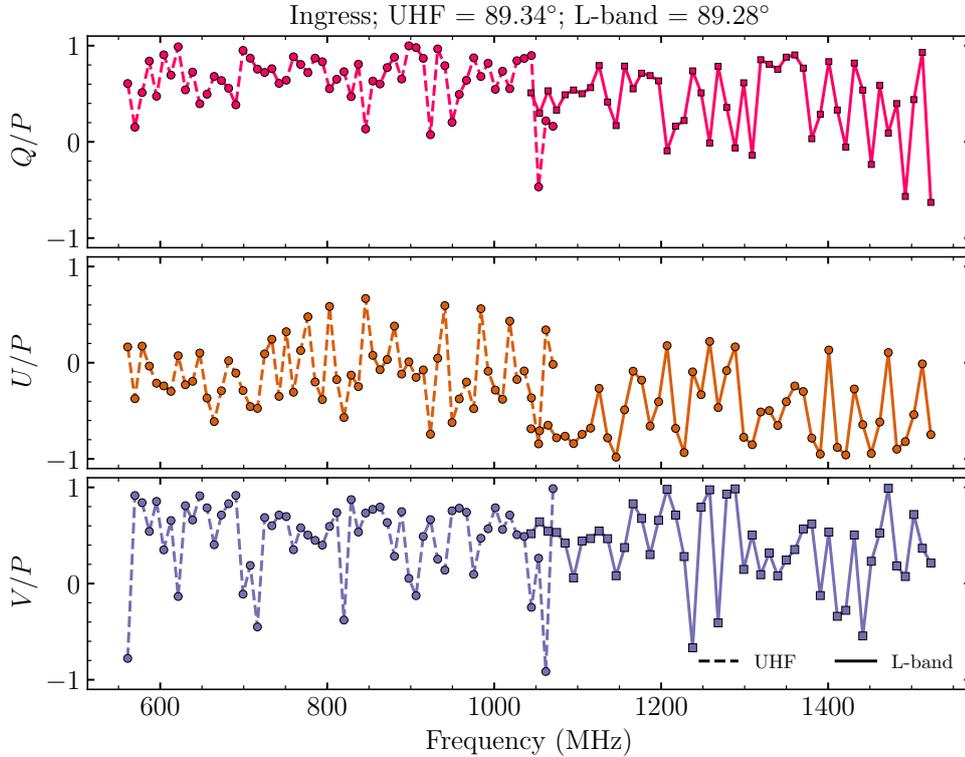


Figure 7.8 Normalised Stokes spectra at the circular polarisation peak within the eclipse ingress phase. All three Stokes spectra are largely uniform across both the L-band and UHF frequency bands. Fluctuations in the Stokes spectra likely originate from a combination of averaging over the scintillation of pulsar A, S/N limitations, and excess RFI that was not fully accounted for.

the circular polarisation comes from the positioning of the RHC and LHC peaks within the eclipse region. The increases in RHC polarisation are localised to the ‘partial transparency’ windows that occur within the ingress and egress phases, while the increased LHC is associated with similar partial transparency windows located near superior conjunction. Partial transparency windows occur when our line of sight to pulsar A intersects the outermost regions encompassed by pulsar B’s closed-field lines. Here, the optical depth of the plasma is small enough that only partial partial synchrotron absorption of the radio pulses takes place.

Similar to standard Faraday rotation, the sign of Faraday conversion is directly related to the direction of the intervening magnetic field along the line of sight. Hence, the handedness of the circular polarisation during ingress, superior conjunction and egress can therefore be related to the average direction of B’s magnetic field at these orbital phases.

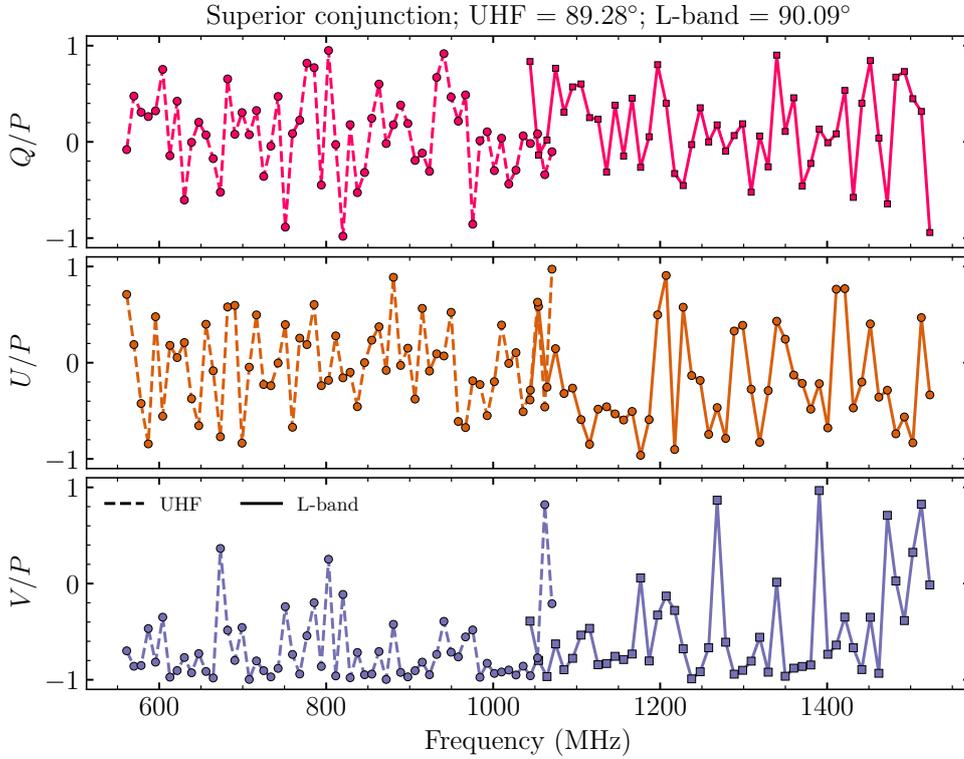


Figure 7.9 Same as Figure 7.8, but for the circular polarisation peak near superior conjunction.

This is illustrated in Figure 7.11, where the approximate line-of-sight to pulsar A is shown for several partial transparency windows detected during an eclipse at UHF frequencies. During ingress, the radiation from pulsar A is partially absorbed when the ‘active’ magnetic pole of pulsar B (the pole from which radio pulses were originally detected, [Noutsos et al. 2020](#)) is maximally pointed in the direction of the Earth. As pulsar A progresses towards superior conjunction, the partially absorbed radiation transitions to intersecting lines of sight where the opposite pole is maximally pointed in our direction. The change in magnetic field direction, therefore, results in the observed flip from RHC to LHC polarisation, reaching a maximum at superior conjunction. Due to the current symmetry in the eclipse shape, this process then reverses as pulsar A moves into the egress phase, with the second peak in RHC occurring at phases where the active pole is again maximally pointed in our direction.

As for the lack of significant linear polarisation being generated within the deepest part of the eclipse, it is possible that it was also converted into circular polarisation

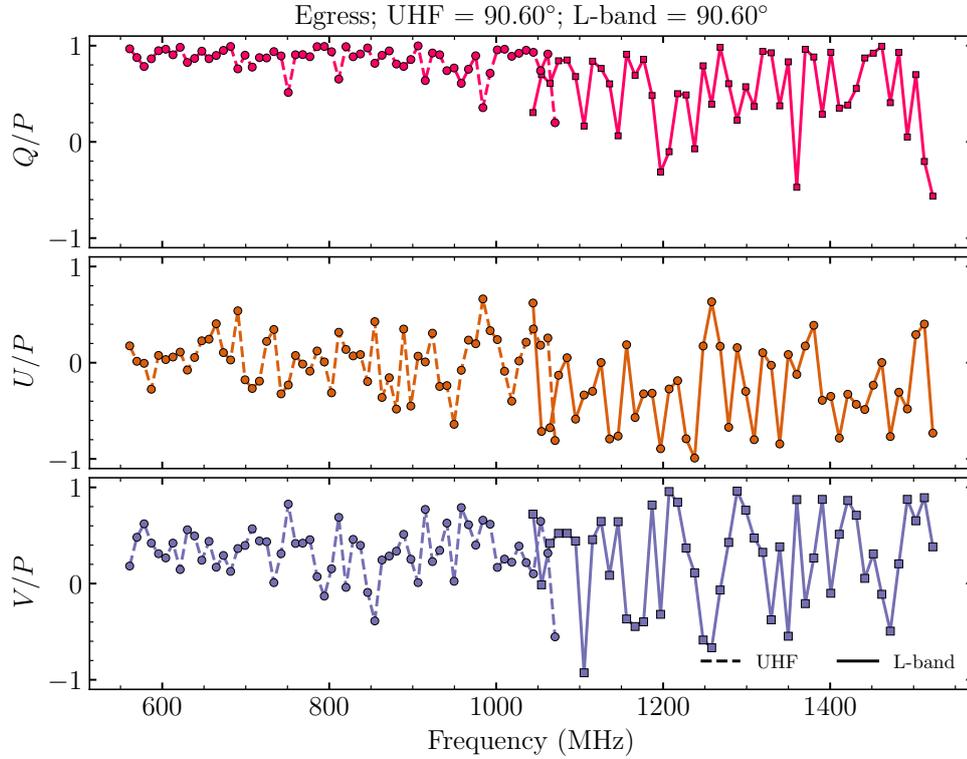


Figure 7.10 Same as Figure 7.8, but for the circular polarisation peak within the eclipse egress phase.

alongside the pre-existing linearly polarised radiation from pulsar A. At the maximum of total polarisation, the total linear component predicted by Lyutikov & Thompson (2005) for a completely unpolarised source makes up only ~ 30 percent of the total intensity for their best-fit eclipse model. The total polarisation at L-band peaks at $\Pi \sim 0.6$, which is consistent with a 30 percent increase above the off-eclipse baseline. At UHF it peaks slightly higher, with $\Pi \sim 0.8$, more than 40 percent higher than the off-eclipse level. The eclipse is somewhat deeper at UHF despite the total linear and circular polarisation light curves appearing largely consistent in amplitude when comparing the two observing bands. Hence, the radiation near superior conjunction at UHF has a higher total polarisation content.

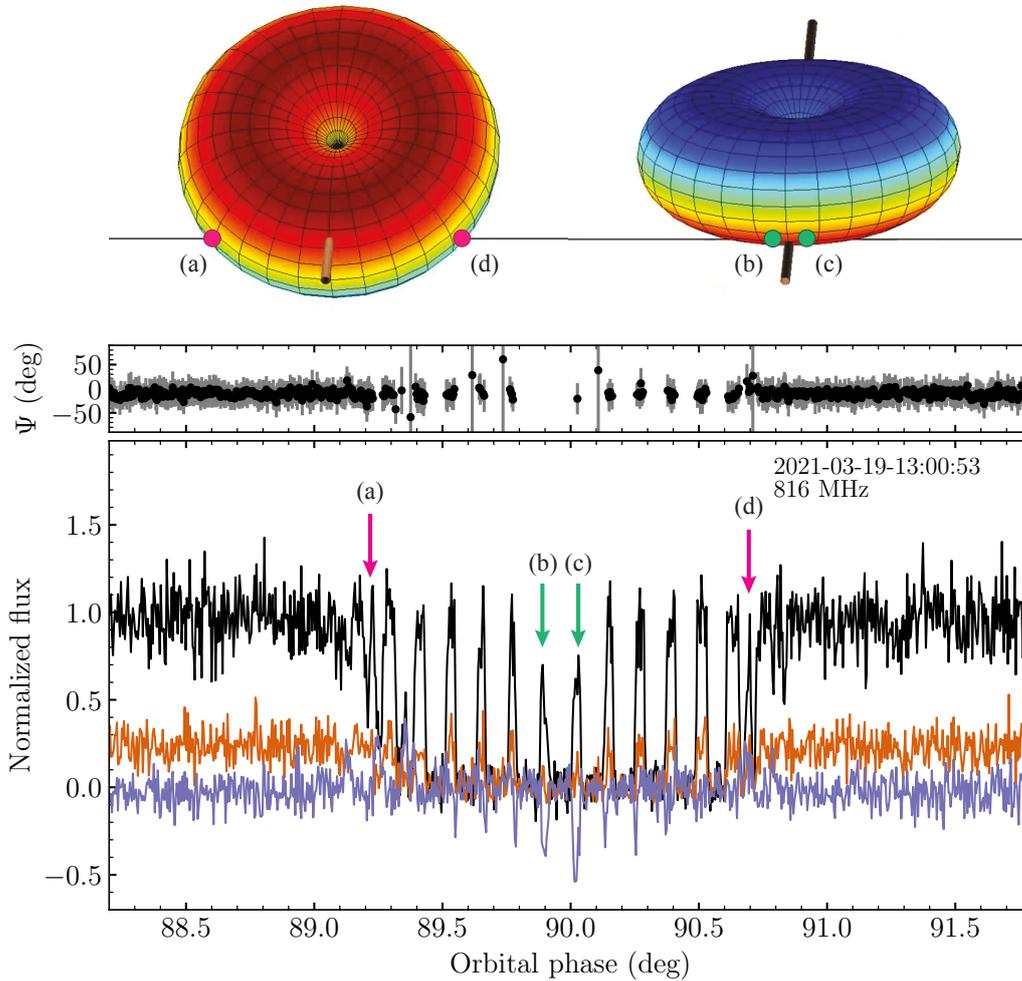


Figure 7.11 Linking birefringent propagation effects in pulsar A to the geometry of pulsar B. Top panel shows the approximate orientation of pulsar B’s toroidal magnetic field during the highlighted partial transparency windows detected in the eclipse light curve shown in the bottom panel. Images of pulsar B’s magnetosphere were adapted from [Breton et al. \(2008\)](#).

7.5 Conclusions

In this chapter we presented the first results from MeerKAT observations of the eclipses detected in the Double Pulsar system PSR J0737–3039A/B. A Bayesian analysis of the eclipse light curve enabled an improved measurement of the geodetic spin-precession rate of pulsar B, where the recovered value of $\Omega_{\text{SO}}^{\text{B}} = 4.98^{+0.39}_{-0.35} \text{ yr}^{-1}$ is consistent with the expected value from GR to within 7.4 percent uncertainty. This enabled an additional update to a theory-independent test of strong-field spin-orbit coupling, where our mea-

surement of $\left(\frac{c^2\sigma_{\text{so}}}{G}\right) = 3.54 \pm 0.27$ is also consistent with GR at the 7.5 percent level. Both measurements represent nearly a factor of two improvement over the results of [Breton et al. \(2008\)](#).

We also performed the first complete polarimetric analysis of the Double Pulsar eclipses using a smaller subset of calibrated MeerKAT observations. The existence of discrete changes in the linear polarisation position angle of pulsar A during the eclipse region and lack of significant variations in rotation measure at L-band and UHF frequencies is somewhat in tension with the claims of [Yuen et al. \(2012\)](#). Our results disfavour the presence of mildly relativistic particles in the magnetotail of pulsar B that would induce a measurable Faraday rotation in the linearly polarised radiation from pulsar A. The birefringent nature of the plasma trapped within the closed-field region of pulsar B's magnetosphere was also confirmed through the detection of high-levels of circular polarisation during portions of the eclipse dominated by partial transparency windows. Modelling of the polarised spectra at the circular polarisation peaks indicates there is little, if any, variation between the Stokes parameters as a function of frequency. This behaviour was ascribed to smearing of the spectra due to pulsar B's rotation within our finite phase bins, averaging over slight differences in individual spectra and the possible presence of multiple birefringence-inducing phenomena. We speculate the sign-changes from RHC to LHC (and vice-versa) are related to changes to the average projected magnetic field direction of pulsar B during the eclipse ingress, superior conjunction and egress.

Improvements to the eclipse model to better account for frequency evolution across the wide fractional bandwidth of MeerKAT may reduce the systematic offsets such that effective ToAs may be produced for pulsar B. This would enable updated timing measurements of the Double Pulsar independent to pulsar A, as well as phase-coherent searches for radio pulsations from pulsar B as we approach its predicted return in 2024 ([Breton, 2009](#)). Phase-coherent timing of pulsar B would also open up the prospect of ascribing particular polarimetric variations in the radiation from pulsar A to specific lines of sight through the closed magnetic field of pulsar B. A more complete picture could also be obtained through including the full radiative transfer of all four Stokes parameters into the eclipse light-curve modelling. However, sampling the high-dimensional parameter space may prove difficult with the current Bayesian inference framework.

The near-future integration of MeerKAT into the full SKA-Mid-1 telescope will provide substantial increases in telescope gain, which can be leveraged to probe the eclipses in far greater detail than what is presented here. SKA observations would not only enable further improved measurements of the pulsar B's geometry via eclipse light-curve modelling, but

the increased sensitivity would allow the birefringent nature of the magnetospheric plasma to be probed at even higher time and frequency resolution than is possible with the current generation of radio telescopes.

8

Conclusions and future prospects

In this thesis we have presented detailed modelling and observational studies of the timing and spectropolarimetric properties of radio pulsars. This included developing an improved understanding of how timing noise varies in non-recycled pulsars, studying the effects glitches have on young pulsar rotational evolution, the spectropolarimetric properties of a recently discovered radio magnetar, and finally, an analysis of high-sensitivity observations of the eclipses detected in the Double Pulsar. Here, we summarise the main findings of these studies and provide comment on potential future measurements that could be made and what development is required to realise them.

8.1 Glitches, timing noise and long-term rotational evolution

In Chapter 3, we expanded the use of existing Bayesian pulsar timing frameworks to a large sample of 300 pulsars regularly observed by the Molonglo Telescope to develop a clearer picture of how long-term timing noise processes vary across the pulsar population. The timing noise of these pulsars was modelled as a red power-law process, which was simultaneously fit for alongside the deterministic pulsar spin-frequency and spin-down rate via the Bayesian pulsar timing package TEMPO_{NEST} (Lentati et al., 2014). Bayesian model selection was employed to distinguish which pulsars showed significant evidence of timing noise in their pulse times of arrival (ToAs). Our initial, model-dependent population analysis revealed a moderate anti-correlation between timing noise and pulsar characteristic ages and a weaker positive correlation with spin-down rate. The second, model independent technique obtained a timing noise scaling relation with pulsar spin-frequency and spin-down to that measured by Shannon & Cordes (2010) despite our comparatively small sample size (280 non-recycled pulsars versus 400) and limited observing baseline (maximum ~ 5 yr versus 40 yr).

Chapter 4 detailed an investigation into the effects of pulsar glitches on the observed long-term rotational evolution of a sample of 74 young pulsars monitored by the Parkes telescope over the last 5-30 yrs. We detected 124 glitches in 52 of these pulsars, of which 74 had not been previously published. Marginalising over the effects of the glitches, we measured the braking indices (n) of 33 pulsars, most of which were found to have $n > 10$. A common origin could not be ruled out between our measurements and similarly large values of n found in predominately non-glitching pulsars by [Parthasarathy et al. \(2020\)](#). Additionally, the rotational evolution of ‘Vela-like’ pulsars (those that had undergone multiple large glitches with significant step-changes in their spin-down) were often dominated by inter-glitch periods of spin-down with $n > 10$, that are interrupted by sudden step-changes in the spin-down rate at each glitch. As a result, the glitch-averaged long-term n of these pulsars is small. We recovered a near one-to-one relationship between the large inter-glitch n and the step-change in spin-down divided by the inter-glitch time interval in the Parkes pulsars that exhibit step-changes in their long-term evolution. This is consistent with theoretical models in which the creeping of superfluid vortices between discrete pinning sites induces an apparent linear glitch recovery process ([Alpar et al., 1984a,b, 1993](#); [Alpar & Baykal, 2006](#)).

The pulsar timing programmes at Molonglo and Parkes are ongoing. Below we outline some of the potential research that may be undertaken with both current and future data collected at the two observatories.

- *Improved population studies of timing noise with astrophysical models*

While the results from both population analysis techniques used in Chapter 3 imply older pulsars that possess lower spin-down rates exhibit less timing noise, the underlying toy-models that were employed make no conclusive statements about the underlying physics responsible for this behaviour. However, it can be easily adapted for use with astrophysically motivated models that make falsifiable predictions on how pulsar timing noise should relate to the physical properties of neutron stars (see, e.g. [Melatos & Link 2014](#)).

- *Searches for common noise processes in young pulsars*

The North American Nanohertz Gravitational-wave Observatory (NANOGrav) recently reported the detection of a common red noise signal among their sample of 45 millisecond pulsars (MSPs; [Arzoumanian et al. 2020](#)). The spectral index and amplitude of the NANOGrav signal is consistent with the expected signal induced by the presence of a stochastic gravitational-wave background emitted by inspiralling

supermassive black hole binaries. The presence of this signal was verified by independent analysis performed on data collected by the Parkes Pulsar Timing Array (Goncharov et al., 2021b). However, the PPTA analysis noted that apparent common signals can be falsely detected if the MSPs in the array share similar (but not identical) intrinsic timing noise processes. Such behaviour could arise if the underlying physics governing timing noise is the same in these pulsars. Since the timing noise detected in MSPs follows the same scaling relation as non-recycled pulsars (Shannon & Cordes, 2010; Parthasarathy et al., 2019), we may expect to find similar ‘common’ signals in young pulsars when the same search techniques are applied. This would provide a means of verifying the gravitational-wave search methods against scenarios in which a truly common signal (i.e one with identical spectral properties) is not expected. Such a detection may provide new insights into similar ‘common’ noise processes that exist in ensemble samples of young pulsars (see Goncharov et al. 2021a on searching for spectral turnovers in a set of MSPs).

- *Can we predict when the next glitch will occur in a pulsar?*

If the vortex-creep model is indeed responsible for the observed behaviour in the Vela-like pulsars, then it may be possible to predict when the next glitch could occur in them. With measurements of the step-change associated with the previous glitch and the post-glitch value of n , the relationship in Section 4.5.2 can be inverted to infer an approximate time until the next glitch (Akbal et al. 2017 showed this method may work for the Vela pulsar). This could enable targeted, multi-wavelength observing campaigns and high-cadence radio observations around the expected glitch epoch. A positive detection (or even a non-detection) of a high-energy transient associated with a glitch could confirm or rule-out a range of glitch models (e.g. Bransgrove et al., 2020). Additional real-time glitch detections would also further our understanding of the internal composition and dynamics of neutron stars (Graber et al., 2018; Ashton et al., 2019a).

- *Measurements of the underlying long-term braking indices of Vela-like pulsars*

Future modifications to our current TEMPONEST-based inference framework may enable us to fit for this small- n process alongside the large- n inter-glitch and timing noise processes in these pulsars. Alternatively, Espinoza et al. (2017) demonstrated a glitch-template matching method that can be used to infer the long-term n of pulsar from local measurements of the spin-down rate. The use of this method could provide an independent check of future measurements of n via extended Bayesian

timing methods.

- *Enhanced timing via profile-domain methods*

Both timing studies in Chapters 3 and 4 relied upon an assumption that the observed pulse profiles remain temporally stable, and could be described by a single, pulsar specific standard profile. However, previous studies have shown the pulse profiles of many pulsars exhibit a range of secular, transient and quasi-periodic variations over time, that can introduce excess timing noise (Lyne et al., 2010; Brook et al., 2016; Goncharov et al., 2021a). Profile-domain timing methods, first introduced by Lentati et al. (2015), avoids biases introduced by variable pulse shapes by directly modelling the observed pulse profile at each epoch while simultaneously fitting the timing model. Demonstrations of the profile-domain method used both real and simulated MSP observations to show it can account for profile stochasticity introduced by short-term jitter, as well as variations introduced by scatter broadening and wide-band profile evolution (Lentati & Shannon, 2015; Lentati et al., 2017a,b). A significant fraction of the young pulsars analysed in this thesis and Parthasarathy et al. (2019, 2020) show evidence of correlated spin-down and profile variations, as well as quasi-periodic emission state switching. Hence the use of profile-domain methods may allow us to break the degeneracy between profile changes and low-significance signals that are currently absorbed by the red-noise modelling of TEMPONEST.

- *Wide-band timing of young pulsars*

The Parkes young-pulsar timing programme has made extensive use of the UWL receiver system since its installation in 2018. The implementation of wideband timing techniques, such as the extended profile-domain method of Lentati et al. (2017a) and the PULSEPORTRAITURE technique developed by Pennucci et al. (2014), are needed to take full advantage of the increased bandwidth and sensitivity afforded by the UWL. In addition to improving the overall quality of recent ToAs, these techniques can also be used to constrain temporal variations in the dispersion and scatter broadening of these pulsars (e.g. Lentati et al., 2017b). Recently PULSEPORTRAITURE was shown to produce consistent ToAs when compared to traditional multi-wavelength timing techniques used by the NANOGrav collaboration (Alam et al., 2021b,a). Hence future iterations of the current data processing and timing pipeline could be used to reprocess the old multi-band data to produce both a simpler and self-consistent set of pulsar timing data.

8.2 Swift J1818.0–1607 and other radio-loud magnetars

Chapter 5 presented the spectropolarimetric analysis of the newly discovered radio-loud magnetar Swift J1818.0–1607 (J1818 hereafter). Using a ~ 3 hr observation of the magnetar taken by the Parkes UWL receiver system, we inferred a rotation measure of $1442.0 \pm 0.2 \text{ rad m}^{-2}$ and found it possessed an unusually steep radio spectrum, with a recovered spectral index of $-2.26_{-0.03}^{+0.02}$. The energy distribution of its single pulses was well described by a log-normal distribution convolved with a Gaussian. Combining measurements of the frequency-dependent dispersion of the single pulses with the aforementioned rotation measure enabled a measurement of the average parallel magnetic field strength along the line of sight, where the value of $\langle B_{\parallel} \rangle \sim 2.5 \mu\text{G}$ is reasonable for its location in the Galaxy. The commonality between the observed spectral and single-pulse properties of J1818 and ‘normal’ rotation-powered pulsars led us to speculate it may represent a transitional pulsar that bridges the two populations. Additionally, the lack of any catalogued supernova remnants at the magnetar’s position hinted at a much older true age than the ~ 410 yr characteristic age reported by [Esposito et al. \(2020\)](#). Drawing comparisons between the early outburst phases of J1818 and PSR J1119–6127 ([Archibald et al., 2016](#)), we suggested the spectrum of J1818 may undergo a similar spectral flattening over the coming months.

In Chapter 6, we showed the latter spectral-flattening prediction indeed came to fruition through the appearance of a new profile component with an inverted radio spectrum. We also identified a set of mode-switching behaviour in the single-pulses detected from the magnetar, including intermittent switching between the original steep-spectrum profile component and the new inverted-spectrum component. Rotating vector model fits to the PA swing at different epochs revealed the magnetic geometry of J1818 is generally stable with time, with almost all observed variations being attributable to variations in the radio emission height. Our detection of a transient, inverted PA swing during one of our observations is difficult to reconcile with standard models of pulsar magnetic fields, particularly since the overall shape of the total intensity profile remained unchanged. We suggested the most likely astrophysical phenomena behind this switching between PA swing directions is the radio emission patch intermittently moves between two, co-located magnetic poles: a more active ‘primary’ pole and a ‘secondary’ pole that is offset by $\sim 30^\circ$ in latitude. Independent evidence for such a magnetic field configuration comes from NICER observations taken by [Hu et al. \(2020b\)](#), where the unusually high pulsed fraction of the detected X-ray profile is best described by either a highly distorted thermal hotspot or a two-component hotspot with differing temperatures.

There are many questions regarding the properties of magnetars that are yet to be answered. For instance, why some are visible as radio pulsars while others are not? Where are they born in the $P-\dot{P}$ diagram? And how are they related to the rotation-powered and high B-field pulsar populations? Below, we describe several future research topics that would extend the analysis of this thesis to both continued observations of Swift J1818.0–1607 and other radio-loud magnetars.

- *Continued monitoring of Swift J1818.0–1607 with Parkes*

The PA-swing flip discussed above was not a one-off, isolated event. Additional flips and reversals have been detected following the publication of Chapter 6, and there is an on-going effort to capture one of these events in action and trigger a set of daily observations following the flip. Tracking of the PA swing as it reverts back to the ‘normal’ configuration via rotating vector model fits would provide a strong test of the intermittent magnetic pole switching hypothesis outlined in the previous section. It would also allow us to build a map of the magnetic field topology connecting the two active regions within the magnetosphere of J1818, a first for any magnetar.

J1818 is also actively being timed by other radio facilities with higher observing cadences than is feasible with Parkes (see e.g. [Champion et al., 2020a](#)), in addition to X-ray satellites in orbit around the Earth ([Hu et al., 2020b](#)). Future collaborations with these groups would allow us to explore whether the timing behaviour of J1818 is affected by the PA-flips. Variations in the surface hotspot temperature or changes in the amount of energy injected by into its surrounding environment associated with the PA-flips may be detectable in the observed X-ray flux from the magnetar.

- *Spectropolarimetry of other radio-loud magnetars*

At least three other radio-loud magnetars have been regularly monitored with the Parkes UWL since its installation in late-2018. Similar studies on the wide-band spectra, polarisation properties, and temporal evolution of their profiles may reveal previously undiscovered pulse phenomenology that could open new insights to the magnetar radio emission mechanism. For instance, [Dai et al. \(2019\)](#) found XTE J1810–197 exhibited significant, correlated variations in the detected linear and circular polarisation across its wide-band spectrum following its 2018 reactivation. This was speculated to be the result of magnetospheric propagation effects. The correlated polarisation variations of this magnetar could be an indicator the radio-pulses emitted by it had traversed a birefringent pair-plasma that is co-rotating with the star. Applying our phenomenological Faraday conversion model (which was used

to analyse the Double Pulsar eclipses in Chapter 7) to this data may allow us to conduct the first direct measurements of the plasma environment that exists within the magnetosphere of a magnetar. A re-analysis of our wide-band monitoring of Swift J1818.0–1607 and archival observations of PSR J1119–6127 performed by Dai et al. (2018) after their respective 2020 and 2016 outbursts may reveal whether similar behaviour is common to recently activated magnetars.

- *Discovering new radio-loud magnetars*

Only six of the thirty known magnetars have been found to emit some form of pulsed radio emission to date (Olausen & Kaspi, 2014)¹. Although theoretical models exist to describe the apparent ‘radio-silent’ nature of most known magnetars (e.g. the ‘fundamental plane’ of magnetar X-ray/spin-down luminosity Rea et al. 2012), it is also possible their pulsed radio emission was either too faint to be detected with sub-100 m class radio telescopes (e.g Lazarus et al., 2012), or they emit highly intermittent radio pulses. Both of these phenomena are seen in the Galactic FRB-analogue, SGR 1935+2154 (CHIME/FRB Collaboration et al., 2020; Bochenek et al., 2020). The detection of sporadic single pulses from this magnetar required dozens to hundreds of hours worth of observing time on both small and large-scale facilities (Kirsten et al., 2020; Zhang et al., 2020), while its transient pulsed emission is both transient in nature and has a significantly lower luminosity compared to other radio magnetars (Zhu et al., 2020). Censuses of magnetars conducted with new extremely sensitive facilities, such as the Five-hundred meter Aperture Spherical Telescope (FAST), MeerKAT and the upcoming SKA-Mid telescope may detect similar emission from the known radio-silent magnetar population. These facilities may also detect new radio-loud magnetars that are either discovered via the detection of high-energy outbursts by all-sky high-energy monitors (e.g. Swift-BAT Barthelmy et al., 2005), or in upcoming blind radio surveys of the Galactic plane, as was the case with PSR J1622–4950 (Levin et al., 2010).

- *Exploring the bridge between magnetars and high-B-field pulsars*

The Thousand-Pulsar Array project on MeerKAT aims to perform at least one observation of every pulsar visible from the Southern Hemisphere (Johnston et al., 2020). A survey of the spectral and polarisation properties of the high B-field pulsars and magnetars observed as part of this programme could unveil links between the radio emission mechanism of the two populations. This could be complemented by

¹<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>

additional wide-band observations of particular objects of interest by Parkes.

8.3 The Double Pulsar

In Chapter 7 we presented an initial analysis of the eclipses in the Double Pulsar system PSR J0737–3039A/B recorded with MeerKAT, the South African SKA-precursor telescope. Using modified variant of the Lyutikov & Thompson (2005) eclipse model and the geometric measurements of Breton et al. (2008), we were able to detect the secular changes induced in the eclipse light-curve of pulsar A by the geodetic spin-precession of pulsar B. Our recovered geodetic precession-rate of $\Omega_{\text{SO}}^{\text{B}} = 4.98^{+0.39}_{-0.35} \text{ yr}^{-1}$ is nearly a factor of two more precise than the previous measurement by Breton et al. (2008). The associated strong-field test of relativistic spin-orbit coupling is consistent with the prediction from general relativity to an uncertainty of 7.5%, which is again nearly a factor of two better than previous measurements.

We also detected significant variations in the pulse phase-averaged position angle of pulsar A after averaging over a set of eclipses detected with the MeerKAT L-band and UHF-band receivers. The non-detection of excess Faraday rotation associated with these PA variations indicates they are likely to be related to differential absorption of the linear polarisation emitted by pulsar A, as opposed to interactions with mildly relativistic particles in the magnetotail of B (Yuen et al., 2012). In addition to the PA variations, we also found the linear polarisation from A undergoes significant Faraday conversion into right-hand circular polarisation within the ingress and egress phases, with sign-flips to (from) left-hand circular polarisation prior to (following) superior conjunction. This significant discovery suggests the magnetospheric plasma of pulsar B is highly birefringent, and may represent a confirmation that it is comprised almost purely of relativistic electron-positron pairs. We speculate the change in handedness may be imparted by the average direction of B’s magnetic field that is probed within the corresponding phases of the eclipse. The lack of any significant frequency-dependence in the circular polarisation is potentially due to the finite bin-width of our observations, which averaged over multiple lines of sight and field reversals within the magnetosphere.

Below we detail a possible improvements that could be made to strengthen the results of our analysis. Additional future extensions to the eclipse modelling that make use of the wide fractional bandwidth of the MeerKAT receivers and long-term monitoring by other facilities are also discussed.

- *Improving the eclipse light-curve model*

Our current implementation of the light-curve model also doesn't take into account the wide fractional bandwidth of the MeerKAT receivers. Averaging over frequency-dependent eclipse evolution may be partially responsible for the observed deviations of the model away from the data. This could be accounted for by modifying the data processing pipeline to produce time-frequency eclipse 'portraits' and extending the light-curve models into the frequency-domain. Additional modifications to how the radial temperature and particle distributions of the plasma are modelled could also improve the light-curve fits by altering how eclipses evolve in width and depth at different frequencies. Indeed a previous study by [Breton et al. \(2012\)](#) indicated the standard eclipse light-curve model of [Lyutikov & Thompson \(2005\)](#) returned poor fits to the eclipse width and depth across a wide range of frequencies.

Additional gains could be made through improved mitigation of eclipse-to-eclipse stochasticity in the magnetosphere of pulsar B. [Breton et al. \(2008\)](#) overcame this issue through a combination of truncating their eclipse data during egress and performing a simultaneous fit to every eclipse in their dataset. As noted in Chapter 7, our current lack of a coherent timing model for pulsar B meant we could not perform the same fit to our current set of eclipses. However, a simultaneous fit may be possible by fitting for the rotation-phase of pulsar B at every epoch. Such an approach would need an additional reference phase parameter for each eclipse epoch (either one per pair of eclipses or one per single eclipse), meaning the current inference framework would need to be adapted to take into account a massive increase in the dimensionality of the problem. This could be solved through the use of Hamiltonian Markov Chain Monte Carlo or diffuse nested sampling algorithms, which are designed to sample large numbers of free parameters. Alternatively, the current framework could be modified to take advantage of recent efforts to implement massive parallelisation into BILBY for accelerated parameter estimation (e.g. [Smith et al., 2020](#)).

- *Timing of B via the eclipses and connecting back to its radio timing*

The aforementioned improvements to the light-curve model may also enable us to extract effective pulse arrival times for pulsar B, as detailed in Section 7.3.1. Phase-coherent timing of pulsar B via the eclipses would greatly enhance our measurements of B's geometry via simultaneous re-fitting of every detected eclipse light curve simultaneously. This would allow us to directly measure the geodetic precession rate (as was done by [Breton et al. 2008](#)), thereby avoiding removing the need for our current hierarchical inference approach. It may also allow us to connect our measurements

back to when pulsar B was detectable via its radio pulses. An eclipse-derived timing solution B would also enable phase-coherent searches for its radio pulses, currently predicted to return between 2024-2030 (Breton, 2009).

- *Enhanced studies of eclipse birefringence*

A timing solution for pulsar B would also enable accurate prediction of its rotation phase throughout individual eclipses. This would reduce or potentially even eliminate the our current process of incoherently average eclipses in orbital phase. Averaging of the eclipses at the rotation-phase of pulsar B (see studies by McLaughlin et al. 2004 and Breton et al. (2008) from when pulsar B was visible as a radio pulsar) would allow us to link the increased levels and sign of right-hand and left-hand circular polarisation detected in the partial transparency windows to the magnetic field direction of pulsar B. Mapping of these variations to specific rotation-phases of pulsar B would result in the first map of the closed magnetic field region of a neutron star.

- *Combining the MeerKAT and Green Bank Telescope eclipse datasets*

Measurements of the precession rate of pulsar B will improve as monitoring of the eclipses by MeerKAT continues. However, a significant improvement could be made if the current measurements could be linked back to the original results of Breton et al. (2008). Monitoring of the Double Pulsar by the 100-m Green Bank Telescope (GBT) has continued to the present day, where observations covering two full orbits of the system are captured every few months (M. McLaughlin private communication). Although the GBT is a factor of $\sim 2-3$ less sensitive than MeerKAT when observing sources at the declination of the Double Pulsar, the long monitoring baseline extends back to the initial discovery of the system in December 2003. In principle, the current data processing and joint eclipse fitting method would require minimal modification to work with the GBT data. If a timing solution becomes available for predicting the rotation phase of pulsar B across the joint MeerKAT-GBT data span, then we could perform a simultaneous fit to every eclipse detected over an 18-year baseline. This would likely return a measurement of pulsar B's geodetic precession rate with a less than 1% uncertainty (as predicted by simulations in Kramer et al. (2021)), making it competitive with current measurements of the five other post-Keplerian parameters inferred from timing pulsar A.

8.4 Final remarks

We have explored the use of Bayesian inference techniques combined with the capabilities of the latest generation of pulsar instrumentation, as well as legacy data sets, to study the magnetospheric and rotational properties of a variety of radio pulsars. This thesis has provided new insights into the impact of rotational irregularities on the long-term evolution of pulsars and the magnetospheric dynamics of a recently activated magnetar. We also used the eclipses of the Double Pulsar to perform a test of gravity and reveal the birefringent nature of the magneto-ionic environment surrounding a neutron star. Ultimately, these results provide a sneak preview of the potential science output of next generation telescopes such as the SKA.

Bibliography

- Abbate, F., Possenti, A., Tiburzi, C., et al. 2020, *Nature Astronomy*
- Abbott, B. P., et al. 2016, *Phys. Rev. Lett.*, 116, 061102
- Abbott, B. P., et al. 2017a, *Phys. Rev. Lett.*, 119, 161101
- Abbott, B. P., et al. 2017b, *ApJ*, 848, L12
- Abbott, B. P., et al. 2018, *Phys. Rev. Lett.*, 121, 161101
- Abdo, A. A., et al. 2013, *ApJS*, 208, 17
- Adhikari, D., et al. 2021, *Phys. Rev. Lett.*, 126, 172502
- Akbal, O., Alpar, M. A., Buchner, S., & Pines, D. 2017, *MNRAS*, 469, 4183
- Alam, M. F., et al. 2021a, *ApJS*, 252, 4
- Alam, M. F., et al. 2021b, *ApJS*, 252, 5
- Alford, M. G., & Schwenzer, K. 2014, *ApJ*, 781, 26
- Alpar, M. A., Anderson, P. W., Pines, D., & Shaham, J. 1984a, *ApJ*, 278, 791
- Alpar, M. A., & Baykal, A. 2006, *MNRAS*, 372, 489
- Alpar, M. A., Chau, H. F., Cheng, K. S., & Pines, D. 1993, *ApJ*, 409, 345
- Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, *Nature*, 300, 728
- Alpar, M. A., Nandkumar, R., & Pines, D. 1985, *ApJ*, 288, 191
- Alpar, M. A., Pines, D., Anderson, P. W., & Shaham, J. 1984b, *ApJ*, 276, 325
- Anderson, P. W., & Itoh, N. 1975, *Nature*, 256, 25
- Antonopoulou, D., Espinoza, C. M., Kuiper, L., & Andersson, N. 2018, *MNRAS*, 473, 1644
- Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, *Science*, 324, 1411
- Archibald, R. F., Kaspi, V. M., Tendulkar, S. P., & Scholz, P. 2016, *ApJ*, 829, L21
- Arons, J. 1993, *ApJ*, 408, 160

- Arons, J., Backer, D. C., Spitkovsky, A., & Kaspi, V. M. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 328, *Binary Radio Pulsars*, ed. F. A. Rasio & I. H. Stairs, 95
- Arzoumanian, Z., Nice, D. J., Taylor, J. H., & Thorsett, S. E. 1994, *ApJ*, 422, 671
- Arzoumanian, Z., et al. 2020, *ApJ*, 905, L34
- Ashton, G., Lasky, P. D., Graber, V., & Palfreyman, J. 2019a, *Nature Astronomy*, 3, 1143
- Ashton, G., Prix, R., & Jones, D. I. 2017, *Phys. Rev. D*, 96, 063004
- Ashton, G., Hübner, M., Lasky, P. D., et al. 2019b, *ApJS*, 241, 27
- Asseo, E., & Khechinashvili, D. 2002, *MNRAS*, 334, 743
- Backer, D. C. 1970a, *Nature*, 228, 1297
- Backer, D. C. 1970b, *Nature*, 228, 42
- Backer, D. C., Rankin, J. M., & Campbell, D. B. 1976, *Nature*, 263, 202
- Bailes, M., Jameson, A., Abbate, F., et al. 2020, *arXiv e-prints*, arXiv:2005.14366
- Bailes, M., Jameson, A., Flynn, C., et al. 2017, *Publ. Astron. Soc. Aust.*, 34, e045
- Balanis, C. A. 2005, *Antenna theory : analysis and design*
- Bannister, K. W., Deller, A. T., Phillips, C., et al. 2019, *Science*, 365, 565
- Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *Space Sci. Rev.*, 120, 143
- Barthelmy, S. D., Gropp, J. D., Kennea, J. A., et al. 2020, *GRB Coordinates Network*, 27696, 1
- Basu, A., Joshi, B. C., Krishnakumar, M. A., et al. 2020, *MNRAS*, 491, 3182
- Baym, G., Pethick, C., Pines, D., & Ruderman, M. 1969, *Nature*, 224, 872
- Beloborodov, A. M. 2009, *ApJ*, 703, 1044
- Beloborodov, A. M., & Thompson, C. 2007, *ApJ*, 657, 967
- Bernardini, F., Perna, R., Gotthelf, E. V., et al. 2011, *MNRAS*, 418, 638

- Bernardini, M. G., D'Avanzo, P., Klingler, N. J., et al. 2020, GRB Coordinates Network, 28055, 1
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *Phys. Rep.*, 203, 1
- Blandford, R. D., & Romani, R. W. 1988, *MNRAS*, 234, 57P
- Blaskiewicz, M., Cordes, J. M., & Wasserman, I. 1991, *ApJ*, 370, 643
- Bochenek, C. D., Ravi, V., Belov, K. V., et al. 2020, *Nature*, 587, 59
- Bock, D. C. J., Large, M. I., & Sadler, E. M. 1999, *AJ*, 117, 1578
- Bonazzola, S., & Gourgoulhon, E. 1996, *A&A*, 312, 675
- Bond, H. E., White, R. L., Becker, R. H., & O'Brien, M. S. 2002, *Publ. Astron. Soc. Pacific*, 114, 1359
- Bransgrove, A., Beloborodov, A. M., & Levin, Y. 2020, *ApJ*, 897, 173
- Breton, R. P. 2009, PhD thesis, McGill University, Canada
- Breton, R. P., Kaspi, V. M., Kramer, M., et al. 2008, *Science*, 321, 104
- Breton, R. P., Kaspi, V. M., McLaughlin, M. A., et al. 2012, *The Astrophysical Journal*, 747, 89
- Brisken, W. F., Macquart, J. P., Gao, J. J., et al. 2010, *ApJ*, 708, 232
- Brook, P. R., Karastergiou, A., Buchner, S., et al. 2014, *ApJ*, 780, L31
- Brook, P. R., Karastergiou, A., Johnston, S., et al. 2016, *MNRAS*, 456, 1374
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, 564, A125
- Buchner, S. 2013, *Astron. Telegram*, 5406
- Buchner, S. J. 2010, *Astron. Telegram*, 2768, 1
- Burgay, M., Esposito, P., Israel, G. L., et al. 2018, in *IAU Symposium, Vol. 337, Pulsar Astrophysics the Next Fifty Years*, ed. P. Weltevrede, B. B. P. Perera, L. L. Preston, & S. Sanidas, 319–321
- Burgay, M., D'Amico, N., Possenti, A., et al. 2003, *Nature*, 426, 531
- Burgay, M., Joshi, B. C., D'Amico, N., et al. 2006, *MNRAS*, 368, 283

- Burke-Spolaor, S., Johnston, S., Bailes, M., et al. 2012, *MNRAS*, 423, 1351
- Caleb, M., Flynn, C., Bailes, M., et al. 2017, *MNRAS*, 468, 3746
- Cameron, P. B., Chandra, P., Ray, A., et al. 2005, *Nature*, 434, 1112
- Camilo, F., Ransom, S. M., Halpern, J. P., & Reynolds, J. 2007a, *ApJ*, 666, L93
- Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2006, *Nature*, 442, 892
- Camilo, F., Reynolds, J., Johnston, S., Halpern, J. P., & Ransom, S. M. 2008, *ApJ*, 679, 681
- Camilo, F., Reynolds, J., Johnston, S., et al. 2007b, *ApJ*, 659, L37
- Camilo, F., Cognard, I., Ransom, S. M., et al. 2007c, *ApJ*, 663, 497
- Camilo, F., Ransom, S. M., Peñalver, J., et al. 2007d, *ApJ*, 669, 561
- Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2016, *ApJ*, 820, 110
- Camilo, F., Scholz, P., Serylak, M., et al. 2018, *ApJ*, 856, 180
- Carlin, J. B., & Melatos, A. 2019, *MNRAS*, 488, 4890
- Carlin, J. B., & Melatos, A. 2020, *MNRAS*, 494, 3383
- Champion, D., Cognard, I., Cruces, M., et al. 2020a, *MNRAS*, 498, 6044
- Champion, D., Desvignes, G., Jankowski, F., et al. 2020b, *The Astronomer's Telegram*, 13559, 1
- Cheng, K. S. 1987a, *ApJ*, 321, 805
- Cheng, K. S. 1987b, *ApJ*, 321, 799
- CHIME/FRB Collaboration, et al. 2020, *Nature*, 587, 54
- CHIME/Pulsar Collaboration, et al. 2021, *ApJS*, 255, 5
- Chukwude, A. E. 2007, , 7, 521
- Chukwude, A. E., & Buchner, S. 2012, *ApJ*, 745, 40
- Coles, W., Hobbs, G., Champion, D. J., Manchester, R. N., & Verbiest, J. P. W. 2011, *MNRAS*, 418, 561

- Contopoulos, I., & Spitkovsky, A. 2006, *ApJ*, 643, 1139
- Cordes, J. M. 1986, *ApJ*, 311, 183
- Cordes, J. M., & Chatterjee, S. 2019, *ARA&A*, 57, 417
- Cordes, J. M., & Downs, G. S. 1985, *ApJS*, 59, 343
- Cordes, J. M., Downs, G. S., & Krause-Polstorff, J. 1988, *ApJ*, 330, 847
- Cordes, J. M., & Helfand, D. J. 1980, *ApJ*, 239, 640
- Cordes, J. M., & Lazio, T. J. W. 2002, arXiv e-prints, astro
- Cordes, J. M., Pidwerbetsky, A., & Lovelace, R. V. E. 1986, *ApJ*, 310, 737
- Cordes, J. M., & Rickett, B. J. 1998, *ApJ*, 507, 846
- Cordes, J. M., & Shannon, R. M. 2008, *ApJ*, 682, 1152
- Cromartie, H. T., et al. 2020, *Nature Astronomy*, 4, 72
- Dai, S., Johnston, S., Weltevrede, P., et al. 2018, *MNRAS*, 480, 3584
- Dai, S., Lower, M. E., Bailes, M., et al. 2019, *ApJ*, 874, L14
- D'Amico, N., Manchester, R. N., Durdin, J. M., & Erickson, W. C. 1985, *Publ. Astron. Soc. Aust.*, 6, 174
- D'Amico, N., Manchester, R. N., Durdin, J. M., et al. 1988, *MNRAS*, 234, 437
- Damour, T., & Ruffini, R. 1974, *Academie des Sciences Paris Comptes Rendus Serie Sciences Mathematiques*, 279, 971
- Damour, T., & Taylor, J. H. 1992, *Phys. Rev. D*, 45, 1840
- De Luca, A. 2017, in *Journal of Physics Conference Series*, Vol. 932, *Journal of Physics Conference Series*, 012006
- de Vries, M., Romani, R. W., Kargaltsev, O., et al. 2021, *ApJ*, 908, 50
- Deller, A. T., Goss, W. M., Brisken, W. F., et al. 2019, *ApJ*, 875, 100
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Nature*, 467, 1081

- Desvignes, G., Kramer, M., Lee, K., et al. 2019, *Science*, 365, 1013
- Detweiler, S. 1979, *ApJ*, 234, 1100
- Dewey, R. J., & Cordes, J. M. 1989, in *NATO Advanced Sci. Institutes (ASI) Series C*, ed. H. Ögelman & E. P. J. van den Heuvel, Vol. 262, 119
- Dexter, J., Deller, A., Bower, G. C., et al. 2017, *MNRAS*, 471, 3563
- Dodson, R., Legge, D., Reynolds, J. E., & McCulloch, P. M. 2003, *ApJ*, 596, 1137
- Dodson, R. G., McCulloch, P. M., & Lewis, D. R. 2002, *ApJ*, 564, L85
- Duncan, R. C., & Thompson, C. 1992, *ApJ*, 392, L9
- Dunn, L., Lower, M. E., & Melatos, A. 2021, *MNRAS*, 504, 3399
- Dunning, A., Bowen, M., Bourne, M., Hayman, D., & Smith, S. L. 2015, in *2015 IEEE-APS Topical Conference on Antennas and Propagation in Wireless Communications (APWC)*, 787–790
- Dyks, J. 2020, *MNRAS*, 495, L118
- Eatough, R. P., et al. 2013, *Nature*, 501, 391
- Edwards, R. T., Bailes, M., van Straten, W., & Britton, M. C. 2001, *MNRAS*, 326, 358
- Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, *MNRAS*, 372, 1549
- Eichler, D., Gedalin, M., & Lyubarsky, Y. 2002, *ApJ*, 578, L121
- Ekers, R. D., Lequeux, J., Moffet, A. T., & Seielstad, G. A. 1969, *ApJ*, 156, L21
- Enoto, T., Sakamoto, T., Younes, G., et al. 2020, *The Astronomer’s Telegram*, 13551, 1
- Espinoza, C. M., Antonopoulou, D., Dodson, R., Stepanova, M., & Scherer, A. 2021, *A&A*, 647, A25
- Espinoza, C. M., Lyne, A. G., Kramer, M., Manchester, R. N., & Kaspi, V. M. 2011a, *ApJ*, 741, L13
- . 2011b, *ApJ*, 741, L13
- Espinoza, C. M., Lyne, A. G., & Stappers, B. W. 2017, *MNRAS*, 466, 147
- Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011c, *MNRAS*, 414, 1679

- Esposito, P., Rea, N., Borghese, A., et al. 2020, ApJ, submitted, arXiv:2004.04083
- Evans, P. A., Gropp, J. D., Kennea, J. A., et al. 2020, GRB Coordinates Network, 27373, 1
- Everett, J. E., & Weisberg, J. M. 2001, ApJ, 553, 341
- Farah, W., Flynn, C., Bailes, M., et al. 2018, MNRAS, 478, 1209
- Farah, W., Flynn, C., Bailes, M., et al. 2019, MNRAS, 1717
- Faucher-Giguère, C.-A., & Kaspi, V. M. 2006, ApJ, 643, 332
- Ferdman, R. D., van Haasteren, R., Bassa, C. G., et al. 2010, Classical and Quantum Gravity, 27, 084014
- Feroz, F., & Hobson, M. P. 2008, MNRAS, 384, 449
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, MNRAS, 398, 1601
- Folkner, W. M., Williams, J. G., Boggs, D. H., Park, R. S., & Kuchynka, P. 2014, Interplanetary Network Progress Report, 42-196, 1
- Fonseca, E., Stairs, I. H., & Thorsett, S. E. 2014, ApJ, 787, 82
- Fowler, L. A., Wright, G. A. E., & Morris, D. 1981, A&A, 93, 54
- Frail, D. A., Kulkarni, S. R., & Bloom, J. S. 1999, Nature, 398, 127
- Frederiks, D. D., Palshin, V. D., Aptekar, R. L., et al. 2007, Astronomy Letters, 33, 19
- Fuentes, J. R., Espinoza, C. M., & Reisenegger, A. 2019, A&A, 630, A115
- Fuentes, J. R., Espinoza, C. M., Reisenegger, A., et al. 2017, A&A, 608, A131
- Gaensler, B. M., Haverkorn, M., Staveley-Smith, L., et al. 2005, Science, 307, 1610
- Gaensler, B. M., Stappers, B. W., Frail, D. A., & Johnston, S. 1998, ApJ, 499, L69
- Gajjar, V., Perez, K., Siemion, A., et al. 2020, The Astronomer's Telegram, 13575, 1
- Gavriil, F. P., Gonzalez, M. E., Gotthelf, E. V., et al. 2008, Science, 319, 1802
- Geyer, M., & Karastergiou, A. 2016, MNRAS, 462, 2587
- Geyer, M., Karastergiou, A., Kondratiev, V. I., et al. 2017, MNRAS, 470, 2659

- Gil, J., Gronkowski, P., & Rudnicki, W. 1984, *A&A*, 132, 312
- Gil, J. A., Jessner, A., Kijak, J., et al. 1994, *A&A*, 282, 45
- Gold, T. 1969, *Nature*, 221, 25
- Goldreich, P. 1970, *ApJ*, 160, L11
- Goldreich, P., & Julian, W. H. 1969, *The Astrophysical Journal*, 157, 869
- Goncharov, B., Reardon, D. J., Shannon, R. M., et al. 2021a, *MNRAS*, 502, 478
- Goncharov, B., Shannon, R. M., Reardon, D. J., et al. 2021b, *ApJ*, 917, L19
- Gotthelf, E. V., Halpern, J. P., Alford, J. A. J., et al. 2019, *ApJ*, 874, L25
- Gögüş, E., Lin, L., Kaneko, Y., et al. 2016, *ApJ*, 829, L25
- Gourgouliatos, K. N., & Cumming, A. 2014, *MNRAS*, 438, 1618
- Graber, V., Cumming, A., & Andersson, N. 2018, *ApJ*, 865, 23
- Granet, C., Bowen, M., Reynolds, J., et al. 2011, in *Proceedings of the 5th European Conference on Antennas and Propagation (EUCAP)*, 918–920
- Granet, C., Zhang, H. Z., Forsyth, A. R., et al. 2005, *IEEE Antennas and Propagation Magazine*, 47, 13
- Green, A. J., Cram, L. E., Large, M. I., & Ye, T. 1999, *ApJS*, 122, 207
- Green, A. J., Reeves, S. N., & Murphy, T. 2014, *Publ. Astron. Soc. Aust.*, 31, e042
- Green, D. A. 2019, *Journal of Astrophysics and Astronomy*, 40, 36
- Greenstein, G. 1970, *Nature*, 227, 791
- Gronwall, C., Gropp, J. D., Kennea, J. A., et al. 2020, *GRB Coordinates Network*, 27746, 1
- Groth, E. J. 1975, *ApJS*, 29, 453
- Gruzinov, A., & Levin, Y. 2019, *ApJ*, 876, 74
- Gunn, J. E., & Ostriker, J. P. 1969, *Nature*, 221, 454
- Gupta, V., Bailes, M., Jameson, A., et al. 2019a, *Astron. Telegram*, 12610

- Gupta, V., Bailes, M., Jameson, A., et al. 2019b, *Astron. Telegram*, 13282
- Gupta, V., Bailes, M., Jameson, A., et al. 2019c, *Astron. Telegram*, 12995
- Gupta, V., Bailes, M., Jameson, A., et al. 2019d, *Astron. Telegram*, 13363
- Halpern, J. P., Gotthelf, E. V., Reynolds, J., Ransom, S. M., & Camilo, F. 2008, *ApJ*, 676, 1178
- Hampson, G., Macleod, A., Beresford, R., et al. 2012, in *2012 International Conference on Electromagnetics in Advanced Applications*, 807–809
- Han, J. L., Manchester, R. N., van Straten, W., & Demorest, P. 2018, *ApJS*, 234, 11
- Hankins, T. H. 1971, *ApJ*, 169, 487
- Hankins, T. H., & Rickett, B. J. 1975, *Methods in Computational Physics*, 14, 55
- Harding, A. K., Contopoulos, I., & Kazanas, D. 1999, *ApJ*, 525, L125
- Harding, A. K., Tademaru, E., & Esposito, L. W. 1978, *ApJ*, 225, 226
- Haskell, B., & Antonopoulou, D. 2014, *MNRAS*, 438, L16
- Haskell, B., Antonopoulou, D., & Barengi, C. 2020, *MNRAS*, 499, 161
- Haskell, B., & Melatos, A. 2015, *International Journal of Modern Physics D*, 24, 1530008
- Hellings, R. W., & Downs, G. S. 1983, *ApJ*, 265, L39
- Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, *ApJ*, 876, L23
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709
- Hinton, S. R. 2016, *The Journal of Open Source Software*, 1, 00045
- Ho, W. C. G. 2011, *MNRAS*, 414, 2567
- Ho, W. C. G., & Andersson, N. 2012, *Nature Physics*, 8, 787
- Ho, W. C. G., Espinoza, C. M., Antonopoulou, D., & Andersson, N. 2015, *Science Advances*, 1, e1500578
- Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, *MNRAS*, 360, 974

- Hobbs, G., Lyne, A. G., Kramer, M., Martin, C. E., & Jordan, C. 2004, MNRAS, 353, 1311
- Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, Classical and Quantum Gravity, 27, 084013
- Hobbs, G., Miller, D., Manchester, R. N., et al. 2011, Publ. Astron. Soc. Aust., 28, 202
- Hobbs, G., Manchester, R. N., Dunning, A., et al. 2020, Publ. Astron. Soc. Aust., 37, e012
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, MNRAS, 369, 655
- Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv e-prints, arXiv:1008.4686
- Hogg, D. W., & Foreman-Mackey, D. 2018, ApJS, 236, 11
- Hotan, A. W., Bailes, M., & Ord, S. M. 2005, MNRAS, 362, 1267
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, Publ. Astron. Soc. Aust., 21, 302
- Howitt, G., Melatos, A., & Delaigle, A. 2018, ApJ, 867, 60
- Hu, C.-P., Strohmayer, T. E., Ray, P. S., et al. 2020a, The Astronomer's Telegram, 13588, 1
- Hu, C.-P., Begiçarslan, B., Güver, T., et al. 2020b, ApJ, 902, 1
- Hulse, R. A., & Taylor, J. H. 1975, ApJ, 195, L51
- Ilie, C. D., Johnston, S., & Weltevrede, P. 2019, MNRAS, 483, 2778
- Ilie, C. D., Weltevrede, P., Johnston, S., & Chen, T. 2020, MNRAS, 491, 3385
- Jackson, J. D. 1998, Classical Electrodynamics, 3rd Edition
- Jacoby, B. A., Bailes, M., Ord, S. M., Edwards, R. T., & Kulkarni, S. R. 2009, ApJ, 699, 2009
- Jankowski, F., van Straten, W., Keane, E. F., et al. 2018, MNRAS, 473, 4436
- Jankowski, F., Bailes, M., Barr, E., et al. 2015a, Astron. Telegram, 8298
- Jankowski, F., Bailes, M., Barr, E., et al. 2015b, Astron. Telegram, 6903
- Jankowski, F., Bailes, M., Barr, E., et al. 2016, Astron. Telegram, 9054

- Jankowski, F., Bailes, M., Barr, E. D., et al. 2017, *Astron. Telegram*, 10770
- Jankowski, F., Bailes, M., van Straten, W., et al. 2019, *MNRAS*, 484, 3691
- Jenet, F. A., Anderson, S. B., Kaspi, V. M., Prince, T. A., & Unwin, S. C. 1998, *The Astrophysical Journal*, 498, 365
- Jenet, F. A., Hobbs, G. B., van Straten, W., et al. 2006, *ApJ*, 653, 1571
- Johnston, S., & Galloway, D. 1999, *MNRAS*, 306, L50
- Johnston, S., & Karastergiou, A. 2017, *MNRAS*, 467, 3493
- Johnston, S., & Karastergiou, A. 2019, *MNRAS*, 485, 640
- Johnston, S., Karastergiou, A., Mitra, D., & Gupta, Y. 2008, *MNRAS*, 388, 261
- Johnston, S., & Kerr, M. 2018, *MNRAS*, 474, 4629
- Johnston, S., & Kramer, M. 2019, *MNRAS*, 490, 4565
- Johnston, S., & Lower, M. E. 2021, *MNRAS* submitted
- Johnston, S., Lyne, A. G., Manchester, R. N., et al. 1992a, *MNRAS*, 255, 401
- Johnston, S., Manchester, R. N., Lyne, A. G., et al. 1992b, *ApJ*, 387, L37
- Johnston, S., & Weisberg, J. M. 2006, *MNRAS*, 368, 1856
- Johnston, S., Karastergiou, A., Keith, M. J., et al. 2020, *MNRAS*, 493, 3608
- Johnston, S., Sobey, C., Dai, S., et al. 2021, *MNRAS*, 502, 1253
- Jonas, J., & the MeerKAT Team. 2016, in *MeerKAT Science: On the Pathway to the SKA*, 1
- Jonas, J. L. 2009, *IEEE Proceedings*, 97, 1522
- Jones, P. B. 1990, *MNRAS*, 246, 364
- Karastergiou, A., Kramer, M., Johnston, S., et al. 2002, *A&A*, 391, 247
- Karuppusamy, R., Desvignes, G., Kramer, M., et al. 2020, *The Astronomer's Telegram*, 13553, 1
- Kaspi, V. M., & Beloborodov, A. M. 2017, *ARA&A*, 55, 261

- Kaspi, V. M., & Helfand, D. J. 2002, in *Astronomical Society of the Pacific Conference Series*, Vol. 271, *Neutron Stars in Supernova Remnants*, ed. P. O. Slane & B. M. Gaensler, 3
- Kaspi, V. M., Manchester, R. N., Johnston, S., Lyne, A. G., & D'Amico, N. 1992, *ApJ*, 399, L155
- Kaspi, V. M., Ransom, S. M., Backer, D. C., et al. 2004, *The Astrophysical Journal*, 613, L137
- Kass, R. E., & Raftery, A. E. 1995, *J. of the American Statistical Association*, 90, 773
- Keane, E. F., & Kramer, M. 2008, *MNRAS*, 391, 2009
- Keane, E. F., Barr, E. D., Jameson, A., et al. 2018, *MNRAS*, 473, 116
- Keith, M. J., Shannon, R. M., & Johnston, S. 2013, *MNRAS*, 432, 3080
- Keith, M. J., Jameson, A., van Straten, W., et al. 2010, *MNRAS*, 409, 619
- Kennett, M., & Melrose, D. 1998, *Publ. Astron. Soc. Aust.*, 15, 211
- Kerr, M. 2019, *The Astronomer's Telegram*, 12481
- Kerr, M., Hobbs, G., Johnston, S., & Shannon, R. M. 2016, *MNRAS*, 455, 1845
- Kerr, M., Johnston, S., Hobbs, G., & Shannon, R. M. 2015, *ApJ*, 809, L11
- Kirsten, F., Snelders, M., Jenkins, M., et al. 2020, *arXiv e-prints*, arXiv:2007.05101
- Kirsten, F., Vlemmings, W., Freire, P., et al. 2014, *A&A*, 565, A43
- Komesaroff, M. M., Ables, J. G., Cooke, D. J., Hamilton, P. A., & McCulloch, P. M. 1973, *Astrophys. Lett.*, 15, 169
- Kramer, M. 1998, *ApJ*, 509, 856
- Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006a, *Science*, 312, 549
- Kramer, M., Stappers, B. W., Jessner, A., Lyne, A. G., & Jordan, C. A. 2007, *MNRAS*, 377, 107
- Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2006b, *Science*, 314, 97

- Kramer, M., Menten, K., Barr, E. D., et al. 2016, in *MeerKAT Science: On the Pathway to the SKA*, 3
- Kramer, M., Stairs, I. H., Venkatraman Krishnan, V., et al. 2021, *MNRAS*, 504, 2094
- Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968a, *Nature*, 220, 340
- Large, M. I., Vaughan, A. E., & Wielebinski, R. 1968b, *Nature*, 220, 753
- Large, M. L., & Vaughan, A. E. 1971, *MNRAS*, 151, 277
- Lattimer, J. M. 2012, *Annual Review of Nuclear and Particle Science*, 62, 485
- Lattimer, J. M., & Prakash, M. 2004, *Science*, 304, 536
- Lazaridis, K., Jessner, A., Kramer, M., et al. 2008, *MNRAS*, 390, 839
- Lazarus, P., Kaspi, V. M., Champion, D. J., Hessels, J. W. T., & Dib, R. 2012, *ApJ*, 744, 97
- Lentati, L., Alexander, P., & Hobson, M. P. 2015, *MNRAS*, 447, 2159
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2014, *MNRAS*, 437, 3004
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, *Phys. Rev. D*, 87, 104021
- Lentati, L., Kerr, M., Dai, S., et al. 2017a, *MNRAS*, 468, 1474
- Lentati, L., & Shannon, R. M. 2015, *MNRAS*, 454, 1058
- Lentati, L., et al. 2016, *MNRAS*, 458, 2161
- Lentati, L., Kerr, M., Dai, S., et al. 2017b, *MNRAS*, 466, 3706
- Levin, L., Bailes, M., Bates, S., et al. 2010, *ApJ*, 721, L33
- Levin, L., Bailes, M., Bates, S. D., et al. 2012, *MNRAS*, 422, 2489
- Levin, L., Lyne, A. G., Desvignes, G., et al. 2019, *MNRAS*, 488, 5251
- Li, X., Levin, Y., & Beloborodov, A. M. 2016, *ApJ*, 833, 189
- Li, X. H., & Han, J. L. 2003, *A&A*, 410, 253
- Lin, J. 1991, *IEEE Transactions on Information Theory*, 37, 145
- Link, B. 2012, *MNRAS*, 422, 1640

- Link, B. 2014, *ApJ*, 789, 141
- Liu, J., Yan, Z., Yuan, J.-P., et al. 2019, *Research in Astronomy and Astrophysics*, 19, 073
- Liu, K., Karuppusamy, R., Cognard, I., et al. 2020, *The Astronomer's Telegram*, 13997, 1
- Livingstone, M. A., Kaspi, V. M., Gavriil, F. P., et al. 2007, *Ap&SS*, 308, 317
- Lorimer, D. R. 2011, *SIGPROC: Pulsar Signal Processing Programs*
- Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy*, Vol. 4
- Lorimer, D. R., & Kramer, M. 2012, *Handbook of Pulsar Astronomy*
- Lorimer, D. R., & Xilouris, K. M. 2000, *ApJ*, 545, 385
- Lorimer, D. R., Faulkner, A. J., Lyne, A. G., et al. 2006, *MNRAS*, 372, 777
- Lower, M. E., Johnston, S., Shannon, R. M., Bailes, M., & Camilo, F. 2021, *MNRAS*, 502, 127
- Lower, M. E., & Shannon, R. M. 2020, *The Astronomer's Telegram*, 13587, 1
- Lower, M. E., Shannon, R. M., Johnston, S., & Bailes, M. 2020a, *ApJ*, 896, L37
- Lower, M. E., Flynn, C., Bailes, M., et al. 2018, *Res. Notes of the AAS*, 2, 139
- Lower, M. E., Bailes, M., Shannon, R. M., et al. 2019, *Res. Notes of the AAS*, 3, 192
- Lower, M. E., Bailes, M., Shannon, R. M., et al. 2020b, *MNRAS*, 494, 228
- . 2020c, *MNRAS*, 494, 228
- Lyne, A., Graham-Smith, F., Weltevrede, P., et al. 2013, *Science*, 342, 598
- Lyne, A., Hobbs, G., Kramer, M., Stairs, I., & Stappers, B. 2010, *Science*, 329, 408
- Lyne, A. G., Jordan, C. A., Graham-Smith, F., et al. 2015, *MNRAS*, 446, 857
- Lyne, A. G., Manchester, R. N., & Taylor, J. H. 1985, *MNRAS*, 213, 613
- Lyne, A. G., Pritchard, R. S., Graham-Smith, F., & Camilo, F. 1996, *Nature*, 381, 497
- Lyne, A. G., Manchester, R. N., Lorimer, D. R., et al. 1998, *MNRAS*, 295, 743
- Lyne, A. G., Burgay, M., Kramer, M., et al. 2004, *Science*, 303, 1153

- Lyutikov, M. 2002, *ApJ*, 580, L65
- Lyutikov, M. 2005, arXiv e-prints, astro
- Lyutikov, M., & Popov, S. 2020, arXiv e-prints, arXiv:2005.05093
- Lyutikov, M., & Thompson, C. 2005, *The Astrophysical Journal*, 634, 1223
- Maan, Y., Joshi, B. C., Surnis, M. P., Bagchi, M., & Manoharan, P. K. 2019, *ApJ*, 882, L9
- Maan, Y., & van Leeuwen, J. 2020, *The Astronomer's Telegram*, 13560, 1
- Majid, W. A., Pearlman, A. B., Dobрева, T., et al. 2017, *ApJ*, 834, L2
- Majid, W. A., Pearlman, A. B., Prince, T. A., Naudet, C. J., & Bansal, K. 2020a, *The Astronomer's Telegram*, 13898, 1
- Majid, W. A., Pearlman, A. B., Prince, T. A., et al. 2020b, *The Astronomer's Telegram*, 13649, 1
- . 2020c, *The Astronomer's Telegram*, 13649, 1
- Manchester, R. 1984, in *Birth and Evolution of Neutron Stars: Issues Raised by Millisecond Pulsars*, ed. S. P. Reynolds & D. R. Stinebring, 227
- Manchester, R. N. 1975, *Publ. Astron. Soc. Aust.*, 2, 334
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, 129, 1993
- Manchester, R. N., Lyne, A. G., Taylor, J. H., et al. 1978, *MNRAS*, 185, 409
- Manchester, R. N., & Taylor, J. H. 1981, *AJ*, 86, 1953
- Manchester, R. N., Lyne, A. G., D'Amico, N., et al. 1996, *MNRAS*, 279, 1235
- Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, *MNRAS*, 328, 17
- Manchester, R. N., Hobbs, G., Bailes, M., et al. 2013, *Publ. Astron. Soc. Aust.*, 30, e017
- Mao, S. A., Gaensler, B. M., Stanimirović, S., et al. 2008, *ApJ*, 688, 1029
- Marcote, B., Nimmo, K., Hessels, J. W. T., et al. 2020, *Nature*, 577, 190
- Margalit, B., Berger, E., & Metzger, B. D. 2019, *ApJ*, 886, 110

- Matsakis, D. N., Taylor, J. H., & Eubanks, T. M. 1997, *A&A*, 326, 924
- Mauch, T., Murphy, T., Buttery, H. J., et al. 2003, *MNRAS*, 342, 1117
- Mazets, E. P., Aptekar, R. L., Cline, T. L., et al. 2008, *ApJ*, 680, 545
- McKinnon, M. M. 1997, *ApJ*, 475, 763
- McLaughlin, M. A. 2013, *Classical and Quantum Gravity*, 30, 224008
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2004, *The Astrophysical Journal*, 616, L131
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, *Nature*, 439, 817
- Melatos, A. 1997, *MNRAS*, 288, 1049
- Melatos, A. 2000, *MNRAS*, 313, 217
- Melatos, A., Dunn, L. M., Suvorova, S., Moran, W., & Evans, R. J. 2020, *ApJ*, 896, 78
- Melatos, A., Howitt, G., & Fulgenzi, W. 2018, *ApJ*, 863, 196
- Melatos, A., & Link, B. 2014, *MNRAS*, 437, 21
- Melatos, A., & Peralta, C. 2007, *ApJ*, 662, L99
- Melatos, A., Peralta, C., & Wyithe, J. S. B. 2008, *ApJ*, 672, 1103
- Melrose, D. B. 2010, *ApJ*, 725, 1600
- Melrose, D. B., & Rafat, M. Z. 2017, in *Journal of Physics Conference Series*, Vol. 932, *Journal of Physics Conference Series*, 012011
- Mereghetti, S., Tiengo, A., Esposito, P., et al. 2005, *ApJ*, 628, 938
- Metzger, B. D., Margalit, B., & Sironi, L. 2019, *MNRAS*, 485, 4091
- Michel, F. C., & Tucker, W. H. 1969, *Nature*, 223, 277
- Miller, M. C., et al. 2019, *ApJ*, 887, L24
- Mitra, D., & Li, X. H. 2004, *A&A*, 421, 215
- Montoli, A., Antonelli, M., Magistrelli, F., & Pizzochero, P. M. 2020, *A&A*, 642, A223
- Morello, V., Barr, E. D., Cooper, S., et al. 2018, *MNRAS*

- Morris, D., Graham, D. A., & Bartel, N. 1981, MNRAS, 194, 7P
- Murphy, T., Mauch, T., Green, A., et al. 2007, MNRAS, 382, 382
- Muslimov, A., & Page, D. 1995, ApJ, 440, L77
- Namkham, N., Jaroenjittichai, P., & Johnston, S. 2019, MNRAS, 487, 5854
- Ng, C. 2018, in IAU Symposium, Vol. 337, Pulsar Astrophysics the Next Fifty Years, 179–182
- Noerdlinger, P. D. 1978, Phys. Rev. Lett., 41, 135
- Noutsos, A., Karastergiou, A., Kramer, M., Johnston, S., & Stappers, B. W. 2009, MNRAS, 396, 1559
- Noutsos, A., Schnitzeler, D. H. F. M., Keane, E. F., Kramer, M., & Johnston, S. 2013, MNRAS, 430, 2281
- Noutsos, A., Desvignes, G., Kramer, M., et al. 2020, A&A, 643, A143
- Olausen, S. A., & Kaspi, V. M. 2014, ApJS, 212, 6
- Ord, S. M., Bailes, M., & van Straten, W. 2002, ApJ, 574, L75
- Ośłowski, S., van Straten, W., Hobbs, G. B., Bailes, M., & Demorest, P. 2011, MNRAS, 418, 1258
- Oswald, L. S., Karastergiou, A., Posselt, B., et al. 2021, MNRAS, 504, 1115
- Owen, B. J., Lindblom, L., Cutler, C., et al. 1998, Phys. Rev. D, 58, 084020
- Özel, F., & Freire, P. 2016, ARA&A, 54, 401
- Pacholczyk, A. G., & Swihart, T. L. 1970, ApJ, 161, 415
- Pacini, F. 1968, Nature, 219, 145
- Paczynski, B. 1992, , 42, 145
- Palfreyman, J., Dickey, J. M., Hotan, A., Ellingsen, S., & van Straten, W. 2018, Nature, 556, 219
- Palfreyman, J. L., Dickey, J. M., Ellingsen, S. P., Jones, I. R., & Hotan, A. W. 2016, ApJ, 820, 64

- Pang, P. T. H., et al. 2021, *ApJ*, 922, 14
- Parthasarathy, A., Shannon, R. M., Johnston, S., et al. 2019, *MNRAS*, 489, 3810
- Parthasarathy, A., Johnston, S., Shannon, R. M., et al. 2020, *MNRAS*, 494, 2012
- Parthasarathy, A., Bailes, M., Shannon, R. M., et al. 2021, *MNRAS*, 502, 407
- Pearlman, A. B., Majid, W. A., Horiuchi, S., et al. 2016, *The Astronomer's Telegram*, 9870, 1
- Pearlman, A. B., Majid, W. A., Prince, T. A., Kocz, J., & Horiuchi, S. 2018, *ApJ*, 866, 160
- Pearlman, A. B., Majid, W. A., Prince, T. A., et al. 2020, *The Astronomer's Telegram*, 13966, 1
- Pennucci, T. T., Demorest, P. B., & Ransom, S. M. 2014, *ApJ*, 790, 93
- Pennucci, T. T., Possenti, A., Esposito, P., et al. 2015, *ApJ*, 808, 81
- Perera, B. B. P., McLaughlin, M. A., Kramer, M., et al. 2010, *ApJ*, 721, 1193
- Perna, R., & Gotthelf, E. V. 2008, *ApJ*, 681, 522
- Perna, R., & Pons, J. A. 2011, *ApJ*, 727, L51
- Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2019, *A&A Rev.*, 27, 4
- Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2021, arXiv e-prints, arXiv:2107.10113
- Petroff, E., Keith, M. J., Johnston, S., van Straten, W., & Shannon, R. M. 2013, *MNRAS*, 435, 1610
- Petrova, S. A. 2001, *A&A*, 378, 883
- Pines, D. 1974, *Nature*, 248, 483
- Pitkin, M. 2018, *J. Open Source Software*, 3, 538
- Pizzochero, P. M., Montoli, A., & Antonelli, M. 2020, *A&A*, 636, A101
- Raaijmakers, G., et al. 2021, *ApJ*, 918, L29
- Radhakrishnan, V., & Cooke, D. J. 1969, *Astrophys. Lett.*, 3, 225

- Radhakrishnan, V., Cooke, D. J., Komesaroff, M. M., & Morris, D. 1969, *Nature*, 221, 443
- Radhakrishnan, V., & Manchester, R. N. 1969, *Nature*, 222, 228
- Radhakrishnan, V., & Srinivasan, G. 1982, *Current Science*, 51, 1096
- Rankin, J. M. 1990, *ApJ*, 352, 247
- Ransom, S. M. 2001, in *Bulletin of the American Astronomical Society*, Vol. 33, American Astronomical Society Meeting Abstracts, 1484
- Rea, N., Pons, J. A., Torres, D. F., & Turolla, R. 2012, *ApJ*, 748, L12
- Reardon, D. J., Coles, W. A., Hobbs, G., et al. 2019, *MNRAS*, 485, 4389
- Reed, B. T., Fattoyev, F. J., Horowitz, C. J., & Piekarewicz, J. 2021, *Phys. Rev. Lett.*, 126, 172503
- Riley, T. E., et al. 2019, *ApJ*, 887, L21
- Riley, T. E., et al. 2021, *ApJ*, 918, L27
- Romani, R. W. 1996, *ApJ*, 470, 469
- Ruderman, M. 1969, *Nature*, 223, 597
- Ruderman, M. A., & Sutherland, P. G. 1975, *ApJ*, 196, 51
- Sarkissian, J., Hobbs, G., Reynolds, J., Palfreyman, J., & Olney, S. 2019, *The Astronomer's Telegram*, 12466, 1
- Sarkissian, J. M., Reynolds, J. E., Hobbs, G., & Harvey-Smith, L. 2017, *Publ. Astron. Soc. Aust.*, 34, e027
- Sazonov, V. N. 1969, *Soviet Ast.*, 13, 396
- Scholz, P., Camilo, F., Sarkissian, J., et al. 2017, *ApJ*, 841, 126
- Serylak, M., Stappers, B. W., Weltevrede, P., et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 295
- Serylak, M., Johnston, S., Kramer, M., et al. 2021, *MNRAS*, 505, 4483
- Shabanova, T. V. 2007, *Ap&SS*, 308, 591
- Shaham, J. 1977, *ApJ*, 214, 251

- Shannon, R. M., & Cordes, J. M. 2010, *ApJ*, 725, 1607
- Shannon, R. M., & Johnston, S. 2013, *MNRAS Letters*, 435, L29
- Shannon, R. M., Lentati, L. T., Kerr, M., et al. 2016, *MNRAS*, 459, 3104
- Shapiro, I. I. 1964, *Phys. Rev. Lett.*, 13, 789
- Shapiro, S. L., & Teukolsky, S. A. 1983, *Black holes, white dwarfs, and neutron stars : the physics of compact objects* (Wiley Interscience, New York)
- Shternin, P., Kirichenko, A., Zyuzin, D., et al. 2019, *ApJ*, 877, 78
- Siegmán, B. C., Manchester, R. N., & Durdin, J. M. 1993, *MNRAS*, 262, 449
- Skilling, J. 2004, *AIP Conf. Proc.*, 735, 395
- Smith, D. A., Guillemot, L., Camilo, F., et al. 2008, *A&A*, 492, 923
- Smith, E., Rankin, J., & Mitra, D. 2013, *MNRAS*, 435, 1984
- Smith, F. G. 1968a, *Nature*, 220, 891
- Smith, F. G. 1968b, *Nature*, 218, 325
- Smith, F. G. 1999, in *Pulsar Timing, General Relativity and the Internal Structure of Neutron Stars*, ed. Z. Arzoumanian, F. Van der Hooft, & E. P. J. van den Heuvel, 151
- Smith, R. J. E., Ashton, G., Vajpeyi, A., & Talbot, C. 2020, *MNRAS*, 498, 4492
- Sobey, C., Bilous, A. V., Grießmeier, J. M., et al. 2019, *MNRAS*, 484, 3646
- Speagle, J. S. 2020, *MNRAS*, 493, 3132
- Staelin, D. H., & Reifenstein, Edward C., I. 1968, *Science*, 162, 1481
- Stairs, I. H. 2003, *Living Reviews in Relativity*, 6, 5
- Stairs, I. H., Lyne, A. G., & Shemar, S. L. 2000, *Nature*, 406, 484
- Stappers, B. W., Keane, E. F., Kramer, M., Possenti, A., & Stairs, I. H. 2018, *Philosophical Transactions of the Royal Society of London Series A*, 376, 20170293
- Stappers, B. W., Archibald, A. M., Hessels, J. W. T., et al. 2014, *ApJ*, 790, 39

- Staveley-Smith, L., Wilson, W. E., Bird, T. S., et al. 1996, *Publ. Astron. Soc. Aust.*, 13, 243
- Stinebring, D. R., Smirnova, T. V., Hankins, T. H., et al. 2000, *ApJ*, 539, 300
- Svinkin, D., Frederiks, D., Hurley, K., et al. 2021, *Nature*, 589, 211
- Tan, C. M., Bassa, C. G., Cooper, S., et al. 2018, *ApJ*, 866, 54
- Tauris, T. M., & Manchester, R. N. 1998, *MNRAS*, 298, 625
- Tauris, T. M., & van den Heuvel, E. P. J. 2006, *Formation and evolution of compact stellar X-ray sources*, Vol. 39, 623–665
- Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *ApJ*, 846, 170
- Taylor, J. H. 1992, *Philosophical Transactions of the Royal Society of London Series A*, 341, 117
- Taylor, J. H., & Cordes, J. M. 1993, *ApJ*, 411, 674
- Thompson, C., & Duncan, R. C. 1993, *ApJ*, 408, 194
- Thompson, C., & Duncan, R. C. 1995, *MNRAS*, 275, 255
- Thompson, C., Lyutikov, M., & Kulkarni, S. R. 2002, *ApJ*, 574, 332
- Thorstensen, J. R., & Armstrong, E. 2005, *AJ*, 130, 759
- Thrane, E., & Talbot, C. 2019, *Publ. Astron. Soc. Aust.*, 36, e010
- Timokhin, A. N. 2010, *MNRAS*, 408, L41
- Torne, P., Eatough, R. P., Karuppusamy, R., et al. 2015, *MNRAS*, 451, L50
- Torne, P., Macías-Pérez, J., Ladjelate, B., et al. 2020a, *A&A*, 640, L2
- Torne, P., et al. 2020b, *The Astronomer’s Telegram*, 14001, 1
- Urama, J. O., Link, B., & Weisberg, J. M. 2006, *MNRAS*, 370, L76
- Usov, V. V. 1992, *Nature*, 357, 472
- van Eysden, C. A., & Melatos, A. 2010, *MNRAS*, 409, 1253
- van Eysden, C. A., & Melatos, A. 2012, *Journal of Low Temperature Physics*, 166, 151

- van Haasteren, R., & Levin, Y. 2013, MNRAS, 428, 1147
- van Straten, W., & Bailes, M. 2011, Publ. Astron. Soc. Aust., 28, 1
- van Straten, W., Demorest, P., & Osłowski, S. 2012, Astronomical Research and Technology, 9, 237
- van Straten, W., Manchester, R. N., Johnston, S., & Reynolds, J. E. 2010, Publ. Astron. Soc. Aust., 27, 104
- Venkatraman Krishnan, V., Bailes, M., van Straten, W., et al. 2019, ApJ, 873, L15
- Venkatraman Krishnan, V., Bailes, M., van Straten, W., et al. 2020, Science, 367, 577
- Venkatraman Krishnan, V., Flynn, C., Farah, W., et al. 2020, MNRAS, 108
- Viganò, D., & Pons, J. A. 2012, MNRAS, 425, 2487
- Viganò, D., Rea, N., Pons, J. A., et al. 2013, MNRAS, 434, 123
- Wadiasingh, Z., & Timokhin, A. 2019, ApJ, 879, 4
- Wang, F. Y. H., Ruderman, M., Halpern, J. P., & Zhu, T. 1998, ApJ, 498, 373
- Wang, N., Manchester, R. N., & Johnston, S. 2007, MNRAS, 377, 1383
- Wang, W., Zhang, B., Chen, X., & Xu, R. 2019, The Astrophysical Journal, 875, 84
- Watts, A., Espinoza, C. M., Xu, R., et al. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 43
- Weltevrede, P., Johnston, S., & Espinoza, C. M. 2011, MNRAS, 411, 1917
- Weltevrede, P., Johnston, S., Manchester, R. N., et al. 2010, Publ. Astron. Soc. Aust., 27, 64
- Whiteoak, J. B. Z., & Green, A. J. 1996, A&AS, 118, 329
- Wijnands, R., & van der Klis, M. 1998, Nature, 394, 344
- Williamson, I. P. 1972, MNRAS, 157, 55
- Wolszczan, A., & Frail, D. A. 1992, Nature, 355, 145
- Yan, W. M., Wang, N., Manchester, R. N., Wen, Z. G., & Yuan, J. P. 2018, MNRAS, 476, 3677

- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, 835, 29
- Yonemaru, N., Kuroyanagi, S., Hobbs, G., et al. 2021, *MNRAS*, 501, 701
- Yu, M., Manchester, R. N., Hobbs, G., et al. 2013, *MNRAS*, 429, 688
- Yuan, J. P., Manchester, R. N., Wang, N., et al. 2017, *MNRAS*, 466, 1234
- Yuan, J. P., Wang, N., Manchester, R. N., & Liu, Z. Y. 2010, *MNRAS*, 404, 289
- Yue, Y. L., Xu, R. X., & Zhu, W. W. 2007, *Advances in Space Research*, 40, 1491
- Yuen, R., Manchester, R. N., Burgay, M., et al. 2012, *The Astrophysical Journal*, 752, L32
- Yuen, R., & Melrose, D. B. 2014, *Publ. Astron. Soc. Aust.*, 31, e039
- Zeiger, B. R., Brisken, W. F., Chatterjee, S., & Goss, W. M. 2008, *ApJ*, 674, 271
- Zhang, B. 2020, *Nature*, 587, 45
- Zhang, B., Harding, A. K., & Muslimov, A. G. 2000, *ApJ*, 531, L135
- Zhang, C. F., Jiang, J. C., Men, Y. P., et al. 2020, *The Astronomer's Telegram*, 13699, 1
- Zhu, W., Wang, B., Zhou, D., et al. 2020, *The Astronomer's Telegram*, 14084, 1
- Zou, W. Z., Wang, N., Wang, H. X., et al. 2004, *MNRAS*, 354, 811



UTMOST rotational results

Table A.1 Astrometric and rotational parameters of all pulsars analysed in this work, including the sky-locations in equatorial coordinates, spin frequencies, spin-down and second spin-frequency derivative. The period, position and DM epoch is MJD 57600 for all pulsars. Errors for RAJ and DEC represent the one-sigma uncertainty on the last digit from TEMPO2. Uncertainties on ν and $\dot{\nu}$ represent the 95 percent confidence intervals scaled to the last digit. Flags indicate: PPTA – pulsar is observed as part of the Parkes Pulsar Timing Array project (Manchester et al. 2013), B – pulsar is in a binary.

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	ν (Hz)	$\dot{\nu}$ (10^{-15} s $^{-2}$)	$\ddot{\nu}$ (10^{-24} s $^{-3}$)	N_{ToA}	T (yr)	Flags
J0030+0451	00:30:27.423(9)	+04:51:39.7(3)	205.530699027(+7, -8)	-0.39(±1)	-	58	1.04	-
J0134-2937	01:34:18.6939(2)	-29:37:17.157(5)	7.30131486798(±8)	-4.17767(±2)	(±0.1)	198	3.42	-
J0151-0635	01:51:22.718(4)	-06:35:02.98(1)	0.682750032508(+6, -5)	-0.20585(±2)	(+0.038, -0.2)	196	3.74	-
J0152-1637	01:52:10.854(1)	-16:37:53.63(3)	1.20085114658(+7, -9)	-1.87464(+2, -3)	(+0.064, -0.098)	119	3.29	-
J0206-4028	02:06:01.2931(1)	-40:28:03.616(1)	1.58591388856(±5)	-3.01013(±2)	(+0.1, -0.08)	131	3.35	-
J0255-5304	02:55:56.2939(4)	-53:04:21.250(4)	2.23359632409(±2)	-0.155852(+7, -8)	(+0.038, -0.054)	270	3.75	-
J0348+0432	03:48:43.639(3)	+04:32:11.45(2)	25.5606365903(+4, -5)	-0.158(±2)	(+26.2, -14.7)	41	2.85	B
J0401-7608	04:01:51.75(1)	-76:08:12.95(5)	1.83400700914(+9, -7)	-5.1927(+3, -4)	(+0.9, -0.4)	110	3.15	-
J0418-4154	04:18:03.7748(4)	-41:54:14.42(6)	1.32079643042(±2)	-2.30176(±6)	(+0.13, -0.58)	99	3.38	-
J0437-4715	04:37:15.8961(6)	-47:15:09.1107(3)	173.68794581(+2, -1)	-1.72(+3, -4)	-	230	1.41	PPTA, B
J0450-1248	04:50:08.7903(2)	-12:48:07.088(8)	2.28303085312(+6, -4)	-0.5358(±1)	(+1.7, -0.9)	81	3.14	-
J0452-1759	04:52:34.119(1)	-17:59:23.15(3)	1.82168155657(+5, -4)	-19.0941(+1, -3)	(+0.08, -0.45)	121	3.73	-
J0525+1115	05:25:56.498(1)	+11:15:18.8(1)	2.82137062404(±2)	-0.58697(+8, -7)	(+0.21, -0.58)	59	3.15	-
J0529-6652	05:29:50.90(3)	-66:52:39.9(3)	1.02486651379(+2, -1)	-16.2526(+3, -4)	(+20.7, -2.7)	59	2.24	-
J0533+0402	05:33:25.828(5)	+04:01:59.7(2)	1.03840237842(±2)	-0.17255(+8, -1)	(+0.46, -0.16)	54	3.14	-
J0536-7543	05:36:30.829(4)	-75:43:54.63(2)	0.802660896061(+1, -2)	-0.37076(±5)	(+0.06, -0.066)	189	3.63	-
J0601-0527	06:01:58.9731(8)	-05:27:50.92(2)	2.52544324442(+6, -7)	-8.30641(+3, -2)	(+0.085, -0.053)	208	3.72	-
J0624-0424	06:24:20.025(1)	-04:24:50.56(4)	0.962392542206(±1)	-0.769(±2)	(+0.12, -0.12)	120	3.14	-
J0627+0706	06:27:44.172(4)	+07:06:33.0(2)	2.10134828979(±1)	-131.6248(±5)	(+0.3, -1.5)	111	3.15	-
J0630-2834	06:30:49.35(1)	-28:34:42.1(2)	0.803583722117(+3, -2)	-4.6323(+1, -9)	(+0.16, -0.13)	87	3.65	-
J0646+0905	06:46:31.025(5)	+09:05:49.6(3)	1.10630072318(+7, -9)	-0.9009(±2)	(+1.3, -1.1)	160	3.01	-
J0659+1414	06:59:48.188(5)	+14:14:19.2(4)	2.59788422925(±2)	-370.7966(±9)	$1^{+1.1}_{-0.5}$	162	3.15	-
J0711-6830	07:11:54.1654(1)	-68:30:47.296(1)	182.117234537(+9, -1)	-0.4928(±3)	(+4.0, -2.5)	43	2.89	PPTA
J0729-1836	07:29:32.30(1)	-18:36:42.1(2)	1.96011842607(+2, -1)	-72.8289(+6, -7)	$-2^{+2.4}_{-1.8}$	169	3.16	-
J0737-3039A	07:37:51.24669(2)	-30:39:40.6895(3)	44.0540680812(+7, -6)	-3.4149(±3)	(+0.82, -0.63)	144	3.49	B
J0738-4042	07:38:32.244(3)	-40:42:39.43(4)	2.66723044109(±5)	-9.805(±2)	-3.5 ± 1.2	243	3.66	-
J0742-2822	07:42:48.91(4)	-28:22:44.0(7)	5.996127853(±2)	-604.187(±1)	(+1.0, -35.8)	180	3.46	-
J0758-1528	07:58:29.061(2)	-15:28:08.333(4)	1.46570344504(+3, -4)	-3.4786(+2, -1)	(+0.21, -0.28)	144	3.12	-

Table A.1 (Continued)

PSRJ	RAJ (hh:mm:ss)	DECJ ($^{\circ}$: $'$: $''$)	ν (Hz)	$\dot{\nu}$ (10^{-15} s $^{-2}$)	$\ddot{\nu}$ (10^{-24} s $^{-3}$)	N_{ToA}	T (yr)	Flags
J0809-4753	08:09:43.834(3)	-47:53:54.85(2)	1.82747830432(+6, -4)	-10.2748(± 2)	(+0.07, -0.32)	79	3.12	-
J0820-1350	08:20:26.407(1)	-13:50:56.32(4)	0.807668884066(± 4)	-1.37174(± 2)	(+0.084, -0.09)	51	3.57	-
J0820-4114	08:20:15.46(1)	-41:14:35.2(1)	1.8333635346(± 7)	-0.0664(± 2)	(+0.48, -0.78)	88	3.39	-
J0835-4510	08:35:20.6(2)	-45:10:33(1)	11.18677868(± 2)	-13161(+5.4, -5.3)	(+19949.0, -7024.5)	1420	4.81	-
J0837+0610	08:37:05.6462(1)	+06:10:15.87(6)	0.785068914181(+2, -3)	-4.19046(+1, -8)	(+0.033, -0.016)	78	3.75	-
J0837-4135	08:37:21.1922(4)	-41:35:14.589(4)	1.33044994223(± 5)	-6.26572(± 3)	(+0.038, -0.09)	140	3.64	-
J0840-5332	08:40:33.726(9)	-53:32:35.95(6)	1.38770592261(± 3)	-3.154(± 1)	(+0.09, -0.23)	59	3.14	-
J0842-4851	08:42:05.4443(9)	-48:51:20.6(1)	1.5519428595(± 9)	-23.0246(+7, -8)	(+2.7, -0.1)	53	3.13	-
J0846-3533	08:46:06.0712(4)	-35:33:40.91(6)	0.895978340332(± 1)	-1.28499(± 4)	(+0.13, -0.3)	65	3.3	-
J0855-3331	08:55:38.421(3)	-33:31:38.99(4)	0.788929778283(+1, -2)	-3.93327(+4, -3)	(+0.2, -0.28)	76	3.12	-
J0856-6137	08:56:59.27(1)	-61:37:52.71(8)	1.03894958562(+2, -3)	-1.813(+1, -8)	(+0.33, -0.11)	48	3.39	-
J0904-4246	09:04:59.083(9)	-42:46:13.4(1)	1.03608336991(+3, -2)	-2.0148(+9, -1)	(+0.24, -0.4)	50	3.37	-
J0904-7459	09:04:10.47(3)	-74:59:41.7(1)	1.81965845843(+1, -7)	-1.5278(+4, -3)	(+1.1, -1.0)	44	2.98	-
J0907-5157	09:07:15.901(3)	-51:57:59.36(2)	3.9438751026(+7, -5)	-28.5412(± 3)	(+0.87, -0.12)	121	3.5	-
J0908-1739	09:08:38.227(4)	-17:39:39.9(1)	2.48987780015(± 4)	-4.1492(± 2)	(+0.27, -0.98)	37	3.16	-
J0908-4913	09:08:35.46(1)	-49:13:05.00(1)	9.36601123448(± 6)	-1324.821(+3, -4)	$4_{-1.5}^{+1.3}$	173	3.48	-
J0909-7212	09:09:35.81(3)	-72:12:08.94(1)	0.733734742115(± 4)	-0.1762(± 2)	(+0.15, -0.41)	31	3.15	-
J0922+0638	09:22:13.85(3)	+06:38:19(1)	2.32217901642(± 3)	-73.985(± 1)	(+2.5, -9.6)	68	3.02	-
J0924-5302	09:24:08.722(4)	-53:02:42.6(3)	1.33987540461(± 2)	-63.431(+1, -9)	(+3.9, -4.1)	137	3.2	-
J0924-5814	09:24:30.82(1)	-58:14:05.10(1)	1.35225504146(+4, -3)	-9.0013(± 2)	(+0.48, -0.28)	72	3.39	-
J0934-5249	09:34:28.237(5)	-52:49:27.30(5)	0.692148271001(+1, -8)	-2.22898(± 3)	(+0.07, -0.14)	152	3.16	-
J0942-5552	09:42:14.88(6)	-55:52:55.1(5)	1.50514304056(± 3)	-51.376(± 2)	$8_{-4.3}^{+4.1}$	150	3.69	-
J0942-5657	09:42:54.422(5)	-56:57:43.21(3)	1.23737204376(± 4)	-60.6357(± 2)	(+0.37, -0.09)	90	3.18	-
J0944-1354	09:44:28.967(1)	-13:54:41.88(2)	1.75357327379(+5, -7)	-0.13929(± 2)	(+0.069, -0.118)	61	3.11	-
J0953+0755	09:53:09.3121(2)	+07:55:36.9(1)	3.95154788907(+9, -8)	-3.58768(± 3)	(+0.021, -0.07)	73	3.71	-
J0955-5304	09:55:29.461(1)	-53:04:16.64(1)	1.15992862694(± 4)	-4.74318(± 1)	(+0.16, -0.078)	155	3.26	-
J0959-4809	09:59:26.212(6)	-48:09:47.47(7)	1.49234604(+5, -4)	-0.1887(+1, -9)	(+0.49, -0.27)	78	3.37	-
J1001-5507	10:01:37.85(5)	-55:07:07.8(5)	0.696073036122(+9, -1)	-24.9925(+7, -6)	$1_{-1.6}^{+0.8}$	138	3.7	-
J1003-4747	10:03:21.529(1)	-47:47:01.190(2)	3.25654170022(± 3)	-21.96544(+9, -6)	(+0.29, -0.36)	69	3.13	-
J1012-5857	10:12:48.470(5)	-58:57:48.50(3)	1.21962189426(± 1)	-26.47184(± 4)	(+0.05, -0.23)	193	3.71	-
J1013-5934	10:13:31.848(2)	-59:34:26.63(1)	2.25784124141(± 1)	-2.83682(+3, -4)	(+0.12, -0.18)	132	3.43	-
J1016-5345	10:16:31.135(4)	-53:45:14.26(3)	1.29940016176(± 2)	-3.25316(+6, -4)	(+0.24, -0.8)	102	3.14	-
J1017-5621	10:17:12.831(1)	-56:21:30.517(7)	1.98624477638(± 5)	-12.38704(± 2)	(+0.2, -0.052)	127	3.17	B
J1017-7156	10:17:51.3172(5)	-71:56:41.596(2)	427.621905026(± 2)	-0.415(+1, -8)	(+10.3, -5.6)	51	3.38	B
J1022+1001	10:22:58.3(1)	+10:01:58(5)	60.7794479207(+6, -9)	-0.1564(± 3)	(+0.36, -0.52)	39	3.43	PPTA
J1032-5911	10:32:04.876(1)	-59:11:54.8(1)	2.15418526489(+2, -3)	-8.3414(± 5)	(+4.9, -0.4)	152	2.51	-
J1034-3224	10:34:19.46(1)	-32:24:26.2(2)	0.869118880561(+3, -2)	-0.17372(± 1)	(+0.14, -0.14)	41	3.74	-
J1036-4926	10:36:13.121(7)	-49:26:21.2(1)	1.95936510762(± 1)	-6.3385(± 3)	(+6.1, -7.5)	30	2.36	-
J1041-1942	10:41:36.191(9)	-19:42:13.7(2)	0.721308789858(+2, -3)	-0.4925(± 1)	(+0.12, -0.13)	34	3.41	-
J1042-5521	10:42:00.4853(9)	-55:21:05.793(6)	0.854067790376(± 2)	-4.89999(+6, -5)	(+0.24, -0.14)	101	3.36	-
J1043-6116	10:43:55.222(8)	-61:16:51.29(8)	3.46491578971(+8, -7)	-124.947(± 2)	(+77.6, -108.0)	78	1.47	-
J1045-4509	10:45:50.1794(5)	-45:09:54.106(6)	133.79314947(± 2)	-0.3153(+8, -6)	(+2.0, -1.0)	38	3.31	PPTA, B
J1046-5813	10:46:18.815(2)	-58:13:51.89(2)	2.70688676045(± 1)	-8.4(+4, -3)	(+0.48, -0.5)	193	3.25	-
J1047-6709	10:47:28.285(5)	-67:09:51.61(4)	5.0389844311(± 1)	-42.8418(+4, -3)	(+0.9, -1.2)	39	2.43	-
J1048-5832	10:48:13.1(1)	-58:32:03(1)	8.0824185121(+4, -5)	-6273.49(+2, -3)	$90_{-51.2}^{+41.8}$	232	3.44	-
J1056-6258	10:56:25.53(1)	-62:58:47.7(1)	2.36714106203(± 1)	-20.057(+8, -7)	(+0.9, -1.0)	198	3.69	-
J1057-5226	10:57:59.068(8)	-52:26:56.10(8)	5.0731886204(+2, -3)	-150.205(+2, -1)	(+0.1, -5.6)	126	2.62	-
J1057-7914	10:57:27.7(1)	-79:14:23.6(3)	0.74216802567(+8, -7)	-0.7321(+4, -3)	(+1.4, -0.8)	41	2.84	-
J1059-5742	10:59:00.8886(4)	-57:42:14.55(3)	0.843879990895(+1, -8)	-3.0668(± 3)	(+0.08, -0.31)	178	3.18	-
J1105-6107	11:05:26.2(1)	-61:07:48.0(8)	15.8222513283(+4, -3)	-3966.97(± 1)	(+6.0, -31.4)	145	2.85	-
J1110-5637	11:10:00.3712(6)	-56:37:32.57(4)	1.79129810299(+1, -2)	-6.6125(+8, -4)	(+1.8, -0.2)	130	3.17	-
J1112-6613	11:12:38.414(4)	-66:13:04.663(2)	2.9920963178(+5, -3)	-7.385(+1, -2)	(+0.1, -1.6)	96	2.93	-
J1112-6926	11:12:50.78(1)	-69:26:32.33(6)	1.21878739947(+4, -3)	-4.1912(± 1)	(+0.52, -0.6)	97	3.22	-
J1114-6100	11:14:22.69(5)	-61:00:32.1(3)	1.13525643826(± 2)	-59.3019(+6, -5)	(+4.8, -4.2)	137	2.45	-
J1116-4122	11:16:43.083(4)	-41:22:44.86(8)	1.06026074416(+2, -3)	-8.955(± 1)	(+0.04, -0.27)	47	3.53	-
J1121-5444	11:21:19.23(1)	-54:44:04.90(1)	1.86641502454(± 2)	-9.7309(+7, -8)	(+2.2, -3.8)	117	2.98	-
J1123-6259	11:23:55.53(1)	-62:59:10.92(8)	3.68409189328(± 1)	-71.2863(+3, -4)	(+5.3, -6.1)	70	2.99	-
J1126-6942	11:26:21.66(4)	-69:42:15.8(1)	1.72586751278(+3, -2)	-9.8111(+9, -1)	-0.69_{-2}^{+6}	34	2.06	-
J1133-6250	11:33:51.3(1)	-62:50:51(1)	0.9776360471(+6, -4)	-0.448(+9, -1)	(+132.0, -20.1)	128	1.22	-
J1136+1551	11:36:03.0946(5)	+15:51:15.9(1)	0.841809871701(+6, -2)	-2.64185(+9, -3)	(+0.34, -0.09)	36	3.63	-

Table A.1 (Continued)

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	ν (Hz)	$\dot{\nu}$ (10^{-15} s^{-2})	$\ddot{\nu}$ (10^{-24} s^{-3})	N_{ToA}	T (yr)	Flags
J1136–5525	11:36:02.2354(5)	–55:25:06.843(5)	2.74188009799(+1, –2)	–61.8(+1, –8)	(+0.6, –3.7)	147	3.53	–
J1141–3322	11:41:42.756(2)	–33:22:37.31(5)	3.43091071248(±3)	–5.4774(±1)	(+0.24, –0.14)	51	3.13	–
J1141–6545	11:41:07.0006(6)	–65:45:19.05(3)	2.53871590079(+6, –7)	–27.7621(±4)	0.48_{-3}^{+7}	273	3.55	B
J1146–6030	11:46:07.7152(1)	–60:30:59.622(9)	3.65798554138(±1)	–23.93026(±4)	(+0.53, –0.3)	169	3.35	–
J1157–6224	11:57:15.208(7)	–62:24:50.90(5)	2.49671858326(+9, –7)	–24.5068(+4, –5)	(+0.9, –0.63)	229	3.55	–
J1202–5820	12:02:28.358(6)	–58:20:33.41(5)	2.20846545731(+7, –6)	–10.3828(±3)	(+0.95, –0.6)	134	3.51	–
J1210–5559	12:10:05.98706(2)	–55:59:03.8501(2)	3.57439188715(±5)	–9.26691(+2, –3)	(+0.032, –0.003)	158	3.51	–
J1224–6407	12:24:22.264(2)	–64:07:53.79(1)	4.61934676092(+4, –5)	–105.6992(+3, –2)	(+0.09, –0.92)	367	3.55	–
J1231–6303	12:31:13.0(1)	–63:03:18(1)	0.74006295676(±3)	–0.0723(+8, –7)	(+2.4, –0.8)	79	3.38	–
J1239–6832	12:39:58.96(2)	–68:32:28.94(9)	0.768094857398(±2)	–7.01054(+7, –8)	(+1.1, –0.9)	60	3.15	–
J1243–6423	12:43:17.120(6)	–64:23:23.92(4)	2.57410111798(±5)	–29.8026(±3)	(+0.0, –0.49)	345	3.63	–
J1253–5820	12:53:28.418(2)	–58:20:40.47(2)	3.91392670035(±7)	–32.2492(±3)	(+1.24, –0.01)	204	3.38	–
J1257–1027	12:57:04.7796(9)	–10:27:04.77(3)	1.61993710136(+2, –1)	–0.94879(±4)	(+0.27, –0.39)	39	3.18	–
J1259–6741	12:59:22.64(1)	–67:41:40.27(6)	1.5075450023(+5, –4)	–1.9434(±1)	(+0.8, –1.6)	42	2.54	–
J1305–6455	13:05:23.47(2)	–64:55:28.5(1)	1.74931666048(+1, –2)	–12.3373(+9, –8)	(+1.2, –2.7)	175	3.47	–
J1306–6617	13:06:38.12(1)	–66:17:21.2(1)	2.11404065516(±2)	–26.7181(±9)	(+1.41, –0.41)	125	3.38	–
J1312–5402	13:12:04.708(2)	–54:02:42.5(2)	1.37333511678(+4, –5)	–0.2765(+2, –1)	(+0.7, –1.03)	36	3.18	–
J1312–5516	13:12:53.533(9)	–55:16:47.3(1)	1.1775198318(±3)	–7.9101(±1)	(+0.3, –0.99)	98	3.07	–
J1319–6056	13:19:20.250(7)	–60:56:46.79(6)	3.51675963943(+2, –9)	–18.8878(+3, –7)	(+0.87, –0.14)	199	3.19	–
J1320–5359	13:20:53.932(2)	–53:59:04.967(3)	3.57477758551(±6)	–118.146(±3)	0.36_{-3}^{+1}	129	3.5	–
J1326–5859	13:26:58.219(7)	–58:59:29.29(7)	2.09207813232(+7, –8)	–14.2359(+5, –4)	(+1.4, –1.0)	306	3.63	–
J1326–6408	13:26:32.433(2)	–64:08:43.80(1)	1.26155286486(+9, –6)	–4.91831(±2)	(+0.122, –0.15)	161	2.85	–
J1326–6700	13:26:02.706(4)	–67:00:50.1(3)	1.84156958655(+4, –3)	–18.037(±2)	(+1.3, –4.2)	132	3.54	–
J1327–6222	13:27:17.36(7)	–62:22:44.7(5)	1.8870445541(+4, –3)	–66.926(±2)	2.37_{-5}^{+7}	308	3.63	–
J1327–6301	13:27:07.4320(3)	–63:01:15.51(2)	5.08957797539(±4)	–39.6313(±1)	(+0.49, –0.29)	238	3.41	–
J1328–4357	13:28:06.4198(5)	–43:57:44.50(8)	1.87722052793(±7)	–10.7557(+4, –3)	0.29 ± 2	93	2.99	–
J1338–6204	13:38:09.247(7)	–62:04:18.7(5)	0.80710212469(±2)	–8.9837(+6, –5)	(+2.9, –4.5)	198	2.47	–
J1350–5115	13:50:16.159(2)	–51:15:24.56(3)	3.38180924275(+8, –6)	–8.6634(+2, –1)	(+4.0, –5.1)	95	2.43	–
J1355–5153	13:55:58.692(2)	–51:53:53.95(2)	1.55206115637(±3)	–6.7736(±1)	(+0.26, –0.23)	123	3.19	–
J1356–5521	13:56:50.49(2)	–55:21:15.2(2)	1.97090897627(±2)	–2.8152(+5, –4)	(+14.3, –9.1)	31	2.39	–
J1359–6038	13:59:58.230(9)	–60:38:07.671(7)	7.84261649234(±2)	–389.488(±1)	$-3_{-1.1}^{+3.3}$	429	3.54	–
J1401–6357	14:01:52.45(1)	–63:57:42.0(1)	1.18651362793(+8, –1)	–23.6871(+5, –4)	(+0.4, –1.3)	245	3.56	–
J1413–6307	14:13:31.32(4)	–63:07:34.6(3)	2.5319550816(±1)	–47.9(+4, –3)	$-4_{-11.1}^{+3.3}$	152	2.43	–
J1418–3921	14:18:50.28(1)	–39:21:18.6(2)	0.911737714389(+2, –3)	–0.73841(+8, –6)	(+0.42, –0.39)	57	3.63	–
J1420–5416	14:20:29.11(1)	–54:16:22.7(1)	1.06863614348(±2)	–0.26515(+6, –7)	(+0.22, –0.3)	74	3.08	–
J1424–5822	14:24:32.130(8)	–58:22:55.7(1)	2.7267557238(+1, –9)	–29.262(±2)	(+120.5, –91.9)	188	1.25	–
J1428–5530	14:28:26.240(3)	–55:30:50.06(4)	1.75348642691(±1)	–6.41562(+4, –3)	0.055 ± 5	168	3.66	–
J1430–6623	14:30:40.732(1)	–66:23:05.546(1)	1.2731663027(±6)	–4.50256(+4, –3)	(+0.032, –0.039)	170	3.57	–
J1435–5954	14:35:00.208(1)	–59:54:49.5(1)	2.11418109778(±4)	–6.9189(±1)	(+0.81, –0.52)	254	3.49	–
J1452–6036	14:52:51.80(1)	–60:36:30.00(8)	6.4519415824(±1)	–60.401(+3, –2)	(+105.1, –110.5)	152	1.32	–
J1453–6413	14:53:32.652(1)	–64:13:16.095(9)	5.571424352(+2, –3)	–85.1854(+2, –1)	(+0.18, –0.24)	234	3.56	–
J1456–6843	14:55:59.914(1)	–68:43:39.49(1)	3.79684009011(+5, –6)	–1.42687(+3, –2)	(+0.044, –0.009)	119	4.21	–
J1457–5122	14:57:40.093(8)	–51:22:54.9(1)	0.57198175779(+1, –2)	–1.73305(+5, –7)	(+0.22, –0.24)	39	3.03	–
J1507–4352	15:07:34.175(4)	–43:52:04.05(1)	3.48725495693(±1)	–19.2672(+5, –4)	(+0.8, –1.4)	56	3.17	–
J1507–6640	15:07:48.634(1)	–66:40:57.86(1)	2.81170331276(+8, –9)	–9.1066(±2)	(+0.22, –0.14)	111	2.97	–
J1511–5414	15:11:51.285(3)	–54:14:40.32(6)	4.99041973147(+8, –6)	–12.072(+1, –2)	(+144.8, –15.3)	101	1.25	–
J1512–5759	15:12:43.13(1)	–58:00:00.43(1)	7.77001479211(±5)	–413.71(±2)	(+1.7, –13.0)	177	3.19	–
J1514–4834	15:14:14.563(2)	–48:34:19.97(4)	2.19857371563(±3)	–4.47652(+9, –8)	(+2.2, –0.3)	51	2.55	–
J1522–5829	15:22:42.244(4)	–58:29:02.815(3)	2.52937565301(+7, –9)	–13.1588(+3, –4)	(+0.25, –1.07)	187	3.21	–
J1527–3931	15:27:58.828(9)	–39:31:34.2(2)	0.41363243246(±1)	–3.26122(+4, –5)	(+0.19, –0.15)	36	3.21	–
J1527–5552	15:27:40.734(4)	–55:52:08.352(6)	0.953544682096(+3, –4)	–10.2459(+2, –1)	(+0.25, –0.0)	134	3.21	–
J1528–3146	15:28:34.952(1)	–31:46:06.944(6)	16.4413569253(+2, –1)	–0.068(+5, –6)	(+1.36, –0.57)	25	3.05	B
J1534–5334	15:34:08.2790(1)	–53:34:19.57(2)	0.730523027415(+2, –3)	–0.76251(±7)	(+0.013, –0.016)	231	3.63	–
J1534–5405	15:34:33.59(1)	–54:05:40.5(2)	3.4519643055(+1, –8)	–18.528(+2, –3)	$4_{-4.5}^{+5.5}$	100	2.43	–
J1539–5626	15:39:14.07(1)	–56:26:26.2(1)	4.10852985562(+3, –6)	–81.894(+3, –1)	(+1.5, –7.0)	179	3.13	–
J1542–5034	15:42:45.32(2)	–50:34:03.66(3)	1.6687581689(+9, –1)	–11.208(+4, –3)	(+5.2, –7.3)	59	2.39	–
J1543+0929	15:43:38.82(2)	+09:29:16.4(5)	1.33609682985(±2)	–0.7773(±7)	(+1.7, –1.6)	28	2.75	–
J1544–5308	15:44:59.8294(6)	–53:08:46.953(9)	5.60055059845(+1, –9)	–1.88991(+3, –4)	(+0.17, –0.26)	164	3.58	–
J1549–4848	15:49:21.027(6)	–48:48:36.1(1)	3.46791653545(+9, –1)	–169.702(±2)	(+63.2, –56.4)	44	1.46	–
J1553–5456	15:53:59.61(1)	–54:56:06.25(1)	0.9247724034(+6, –5)	–13.4399(±1)	(+36.6, –44.9)	73	1.21	–

Table A.1 (Continued)

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	ν (Hz)	$\dot{\nu}$ (10^{-15} s $^{-2}$)	$\ddot{\nu}$ (10^{-24} s $^{-3}$)	N_{ToA}	T (yr)	Flags
J1555–3134	15:55:17.947(2)	–31:34:20.3(1)	1.93009273101(±2)	–0.23061(±7)	(+0.34, –0.11)	49	3.19	–
J1557–4258	15:57:00.25445(6)	–42:58:12.35(1)	3.0377858242(+1, –2)	–3.04646(+8, –7)	(+0.05, –0.131)	90	3.54	–
J1559–4438	15:59:41.525(1)	–44:38:45.85(3)	3.89018703598(+2, –3)	–15.4484(±2)	(+0.53, –0.27)	66	3.37	–
J1559–5545	15:59:20.7(1)	–55:45:47(1)	1.04464080328(±5)	–21.733(±3)	(+12.5, –3.9)	87	3.49	–
J1600–3053	16:00:51.8941(7)	–30:53:49.70(3)	277.937706823(±2)	–0.687(±5)	(+47.9, –46.2)	32	3.39	PPTA, B
J1600–5044	16:00:53.033(5)	–50:44:20.93(8)	5.19197591119(±2)	–136.452(+1, –9)	(+3.1, –1.2)	175	3.63	–
J1603–2531	16:03:04.8253(6)	–25:31:47.9(4)	3.53267858106(±4)	–19.8906(±2)	(+1.13, –0.29)	29	3.36	–
J1603–2712	16:03:08.036(1)	–27:13:27.0(7)	1.28482652061(±4)	–4.9683(+1, –2)	(+0.11, –0.28)	29	3.0	–
J1603–7202	16:03:35.6736(9)	–72:02:32.795(7)	67.3765811129(+1, –2)	–0.074(+6, –5)	(+0.66, –0.56)	38	3.11	PPTA, B
J1604–4909	16:04:22.985(2)	–49:09:58.33(5)	3.05419496456(±5)	–9.4894(±3)	(+1.25, –0.25)	113	3.47	–
J1605–5257	16:05:16.265(3)	–52:57:34.80(5)	1.51972586121(±1)	–0.59109(+4, –3)	(+0.12, –0.43)	162	3.75	–
J1613–4714	16:13:29.018(4)	–47:14:26.41(8)	2.61522196138(±2)	–4.33493(±6)	(+0.31, –0.23)	65	3.35	–
J1622–4950	16:22:44.80(3)	–49:50:54.5(5)	0.2311087(±3)	–526(+54.3, –56.9)	(+5145.6, –7084.3)	77	1.21	–
J1623–0908	16:23:17.658(4)	–09:08:48.9(3)	0.783424111867(+2, –1)	–1.58401(±5)	(+0.36, –0.49)	31	3.13	–
J1623–4256	16:23:48.291(6)	–42:56:52.6(1)	2.74279572052(±1)	–7.5624(+8, –7)	(+1.16, –0.06)	57	3.47	–
J1626–4537	16:26:48.94(1)	–45:37:25.8(6)	2.701641249(±2)	–60.541(±4)	(+190.9, –230.7)	35	1.21	–
J1633–4453	16:33:47.03(3)	–44:53:08.58(7)	2.2908895877(±3)	–32.539(+5, –6)	(+327.5, –153.7)	33	1.24	–
J1633–5015	16:33:00.0861(1)	–50:15:08.358(3)	2.83973605453(±1)	–30.54746(±4)	(+0.128, –0.07)	110	3.74	–
J1639–4604	16:39:21.198(3)	–46:04:33.23(7)	1.88992880367(±3)	–20.60947(±9)	(+0.69, –0.88)	57	2.47	–
J1644–4559	16:44:49.234(6)	–45:59:10.3(1)	2.19742452445(+4, –3)	–96.9653(+4, –5)	$1_{-1.4}^{+1.1}$	648	4.1	–
J1646–6831	16:46:54.91(3)	–68:31:51.7(1)	0.560031669373(±2)	–0.5337(±1)	(+0.074, –0.14)	27	3.35	–
J1651–4246	16:51:48.797(6)	–42:46:09.97(1)	1.18472094037(+4, –5)	–6.662(+3, –2)	(+0.72, –0.76)	148	3.46	–
J1651–5222	16:51:42.962(2)	–52:22:58.38(3)	1.574658888(±1)	–4.48968(±3)	(+0.24, –0.017)	95	3.38	–
J1651–5255	16:51:41.41(1)	–52:55:47.7(2)	1.12291858733(±1)	–2.6677(±6)	(+0.94, –0.04)	71	3.14	–
J1652–2404	16:52:58.50(5)	–24:03:54(7)	0.586943472123(+4, –3)	–1.0877(±1)	(+0.29, –0.26)	27	2.87	–
J1700–3312	17:00:52.96(2)	–33:12:45(1)	0.736209097583(+6, –5)	–2.5543(±2)	(+0.19, –0.37)	53	3.36	–
J1701–3726	17:01:18.45(1)	–37:26:27.2(5)	0.407395359535(±2)	–1.84611(+5, –6)	(+0.23, –0.21)	62	3.17	–
J1703–1846	17:03:51.102(9)	–18:46:13(1)	1.24325189271(±3)	–2.67613(±1)	(+0.06, –0.21)	34	3.18	–
J1703–3241	17:03:22.514(2)	–32:41:48.5(1)	0.825228539025(±5)	–0.44787(±2)	(+0.057, –0.067)	81	3.54	–
J1703–4851	17:03:54.541(7)	–48:52:01.04(1)	0.716124628374(±2)	–2.60178(+5, –7)	(+0.22, –0.09)	50	3.36	–
J1705–1906	17:05:36.093(2)	–19:06:39.2(3)	3.34458679304(+6, –7)	–46.2498(±3)	(+0.27, –0.89)	78	3.52	–
J1705–3423	17:05:42.362(3)	–34:23:43.1(2)	3.91501633777(±1)	–16.4861(+6, –5)	(+0.6, –1.6)	112	3.43	–
J1707–4053	17:07:21.78(2)	–40:53:55.1(9)	1.72111797519(±2)	–5.6882(+5, –4)	(+2.0, –1.5)	57	3.51	–
J1708–3426	17:08:57.79(1)	–34:26:44.0(6)	1.44484514046(+3, –5)	–8.7827(±2)	(+0.42, –0.59)	52	3.35	–
J1709–1640	17:09:26.452(4)	–16:40:59.2(4)	1.53125350203(+1, –7)	–14.8003(+5, –6)	0.61 ± 2	38	3.63	–
J1709–4429	17:09:42.62(5)	–44:29:12(1)	9.7542901224(+6, –4)	–8850.16(+6, –8)	$196_{-20.1}^{+31.4}$	111	3.5	–
J1711–5350	17:11:53.13(1)	–53:50:18.3(2)	1.11205916031(±1)	–19.2133(±5)	(+0.47, –0.73)	46	3.1	–
J1715–4034	17:15:40.92(3)	–40:34:18(1)	0.482589475307(+7, –6)	–0.7063(±2)	(+0.48, –0.52)	76	3.5	–
J1717–3425	17:17:20.30(1)	–34:25:00.31(8)	1.52368077976(+5, –6)	–22.72(±2)	$2_{-1.6}^{+1.7}$	56	2.38	–
J1717–4054	17:17:52.31(1)	–41:03:13.0(4)	1.12648202933(+6, –8)	–4.7161(+8, –5)	-1.54_{-2}^{+9}	31	4.09	–
J1720–1633	17:20:25.23(1)	–16:33:35(1)	0.638730665146(±3)	–2.3719(+9, –1)	(+0.23, –0.06)	41	3.13	–
J1720–2933	17:20:34.131(5)	–29:33:14.0(5)	1.61173637049(+3, –2)	–1.9396(±1)	(+0.22, –0.22)	43	3.29	–
J1722–3207	17:22:02.9641(1)	–32:07:45.07(6)	2.09574210095(±3)	–2.8316(±1)	(+0.2, –0.22)	89	3.17	–
J1722–3712	17:22:59.17(5)	–37:12:03.(2)	4.23402576366(±4)	–194.486(±2)	$16_{-12.7}^{+13.7}$	116	3.1	–
J1727–2739	17:27:30.98(3)	–27:38:53(4)	0.7733354277(+3, –2)	–0.6399(+5, –6)	$3_{-1.7}^{+15.8}$	34	2.42	–
J1730–2304	17:30:21.682(4)	–23:04:30(1)	123.110287079(+1, –2)	–0.3023(+2, –4)	(+1.6, –1.2)	42	3.3	PPTA
J1731–4744	17:31:42.21(1)	–47:44:38.7(4)	1.2049313854(+2, –3)	–237.394(±5)	$-9_{-0.6}^{+3.2}$	145	3.58	–
J1733–2228	17:33:26.43(3)	–22:28:37(10)	1.14720621377(±5)	–0.0585(+2, –1)	$23_{-14.1}^{+0.0}$	40	3.05	–
J1736–2457	17:36:45.4(1)	–24:57:50(33)	0.3784689286(+3, –2)	–0.452(+5, –6)	(+174.6, –205.3)	25	1.14	–
J1739–2903	17:39:34.285(2)	–29:03:03.96(3)	3.09704932896(+9, –8)	–75.5355(+3, –4)	(+0.36, –0.6)	88	3.02	–
J1740–3015	17:40:33.98(5)	–30:15:22(5)	1.647450502(+2, –3)	–1263.51(+7, –9)	(+46.0, –51.7)	229	3.47	–
J1741–3927	17:41:18.079(1)	–39:27:38.12(7)	1.95231526583(±1)	–6.4645(+5, –6)	2.81_{-5}^{+3}	74	3.14	–
J1743–3150	17:43:36.710(8)	–31:50:22.7(9)	0.414138298084(±1)	–20.7152(+4, –5)	(+0.13, –0.09)	84	3.16	–
J1745–3040	17:45:56.3081(6)	–30:40:23.30(6)	2.72156341619(+1, –2)	–79.04005(+8, –9)	(+0.23, –0.1)	110	3.5	–
J1751–4657	17:51:42.185(1)	–46:57:26.72(4)	1.34706694407(+4, –3)	–2.35478(+1, –2)	(+0.09, –0.083)	53	3.61	–
J1752–2806	17:52:58.707(8)	–28:06:36(1)	1.77757096075(±6)	–25.6877(+8, –7)	0.51_{-4}^{+3}	145	4.1	–
J1757–2421	17:57:29.37(1)	–24:19:54(10)	4.2715099866(+2, –3)	–236.544(+5, –4)	(+166.8, –147.0)	66	1.31	–
J1759–2205	17:59:24.164(4)	–22:05:33(2)	2.16928428064(+1, –2)	–51.0746(+7, –6)	(+1.52, –0.26)	54	3.02	–
J1759–3107	17:59:22.056(4)	–31:07:21.8(5)	0.926822758345(±3)	–3.24135(±9)	(+0.81, –0.48)	40	2.39	–
J1801–0357	18:01:22.628(3)	–03:57:55.7(2)	1.08519559579(±4)	–3.8928(±1)	(+0.78, –0.46)	29	2.34	–

Table A.1 (Continued)

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	ν (Hz)	$\dot{\nu}$ (10^{-15} s^{-2})	$\ddot{\nu}$ (10^{-24} s^{-3})	N_{ToA}	T (yr)	Flags
J1801–2920	18:01:46.839(3)	–29:20:38.1(3)	0.924290873961(±1)	–2.81266(±3)	(+0.24, –0.02)	60	3.26	–
J1803–2137	18:03:51.4(1)	–21:37:07.(27)	7.478883401(+4, –6)	–7488(+1.1, –0.9)	$283^{+29.1}_{-39.6}$	52	1.29	–
J1805–1504	18:05:06.1(2)	–15:04:36(10)	0.84654711(±1)	–0.31(±2)	(+378.9, –625.7)	28	1.25	–
J1807–0847	18:07:38.0259(2)	–08:47:43.28(1)	6.10771328217(+6, –5)	–1.06808(+2, –1)	(+0.12, –0.057)	74	3.7	–
J1807–2715	18:07:08.4918(3)	–27:15:02.07(5)	1.20804374592(±6)	–17.8128(±2)	(+0.65, –0.2)	77	3.12	–
J1808–0813	18:08:09.432(1)	–08:13:01.8(4)	1.14149384538(±5)	–1.6108(±2)	(+0.13, –0.35)	32	3.38	–
J1809–2109	18:09:14.32(3)	–21:09:02.(5)	1.42365721129(±4)	–7.747(±1)	(+22.7, –13.0)	29	2.4	–
J1810–5338	18:10:44.473(3)	–53:38:07.631(6)	3.8306868647(±4)	–5.6604(±1)	(+0.78, –0.49)	32	2.96	–
J1816–2650	18:16:35.399(6)	–26:49:53(1)	1.68666719259(+2, –3)	–0.18919(+9, –1)	(+0.75, –0.83)	51	3.35	–
J1818–1422	18:18:23.77(1)	–14:22:39(1)	3.43064845763(±2)	–23.9924(+6, –7)	(+1.2, –3.0)	51	3.02	–
J1820–0427	18:20:52.559(2)	–04:27:37.9(1)	1.67201171071(+6, –9)	–17.6967(+5, –3)	-0.64^{+5}_{-6}	55	3.63	–
J1822–2256	18:22:58.95(4)	–22:56:29(16)	0.53354117731(±2)	–0.38531(±5)	(+0.31, –0.14)	57	3.33	–
J1823–0154	18:23:52.138(3)	–01:54:04.94(1)	1.31617369972(+3, –2)	–1.95718(±9)	(+0.41, –0.34)	35	3.14	–
J1823–1115	18:23:40.3(1)	–11:15:11(1)	3.57360247363(±2)	–17.5869(+7, –8)	(+2.8, –2.4)	53	3.19	B
J1823–3106	18:23:46.819(4)	–31:06:48.0(3)	3.52042950493(+1, –9)	–36.3594(+4, –6)	0.69 ± 3	35	3.11	–
J1824–0127	18:24:53.43(1)	–01:27:51.4(4)	0.400084842618(±3)	–0.62531(+8, –1)	(+0.81, –0.82)	30	2.4	–
J1824–1945	18:24:00.4360(4)	–19:45:44.5(8)	5.28154642765(±1)	–146.2029(+8, –7)	$-2^{+1.9}_{-1.9}$	95	3.38	–
J1825–0935	18:25:30.62(6)	–09:35:22(4)	1.3003801253(±1)	–88.397(±3)	(+59.4, –52.1)	144	3.84	–
J1827–0750	18:27:02.7071(6)	–07:50:15.4(2)	3.69682174881(+6, –5)	–21.224(±2)	(+8.0, –6.7)	49	2.36	–
J1829–1751	18:29:43.15(1)	–17:51:13(1)	3.25587939511(+2, –3)	–58.78(±1)	(+0.9, –1.1)	79	3.41	–
J1830–1135	18:30:01.787(6)	–11:35:27(6)	0.160730937421(±6)	–1.2319(±2)	(+0.51, –0.61)	40	2.99	–
J1832–0827	18:32:37.013(2)	–08:27:03.16(1)	1.54478860938(±6)	–152.4959(+3, –2)	-0.39 ± 3	93	3.21	–
J1833–0338	18:33:42.028(8)	–03:39:08.00(3)	1.45617035595(+1, –9)	–88.0832(+4, –5)	$1^{+1.1}_{-1.7}$	102	3.19	–
J1833–0827	18:33:40.245(2)	–08:27:30.9(1)	11.7247184586(±5)	–1261.993(±2)	(+4.1, –5.2)	54	2.96	–
J1834–0426	18:34:25.621(3)	–04:26:15.7(2)	3.44698922332(±3)	–0.8605(±1)	(+1.0, –0.8)	53	3.22	–
J1835–1020	18:35:57.44(3)	–10:19:51(2)	3.3063181115(±1)	–64.651(±5)	(+5.2, –16.0)	53	3.23	–
J1836–0436	18:36:51.77(1)	–04:36:37.3(7)	2.82296346028(±4)	–13.232(±1)	(+11.7, –48.8)	31	2.42	–
J1836–1008	18:36:53.922(3)	–10:08:09.39(2)	1.77708391524(+8, –1)	–37.1805(+8, –4)	(+1.2, –1.3)	74	3.53	–
J1837–0653	18:37:14.53(7)	–06:52:55(5)	0.52471147608(+4, –5)	–0.194(+2, –1)	(+6.1, –6.1)	50	2.24	–
J1840–0809	18:40:33.365(6)	–08:09:03.62(4)	1.04638272527(±3)	–2.57306(+9, –1)	(+0.17, –0.2)	75	3.43	–
J1840–0815	18:40:13.756(9)	–08:15:08.88(4)	0.912041662037(±3)	–2.01835(±8)	(+0.6, –1.3)	51	2.44	–
J1841–0425	18:41:05.663(5)	–04:25:19.5(2)	5.37198570613(±4)	–184.318(±1)	(+3.2, –6.3)	31	2.24	–
J1841+0912	18:41:55.921(7)	+09:12:07.29(2)	2.62246808546(+3, –4)	–7.505(+2, –1)	(+0.3, –4.2)	29	3.23	–
J1842–0359	18:42:26.49(1)	–04:00:01.52(7)	0.543494594895(±5)	–0.1501(±1)	(+0.4, –1.0)	77	3.44	–
J1843–0000	18:43:27.965(9)	–00:00:41.5(2)	1.13593208324(+1, –2)	–10.0305(±7)	(+1.0, –0.7)	60	3.19	–
J1844–0433	18:44:33.446(3)	–04:33:12.5(1)	1.00905187281(±2)	–3.98545(+4, –5)	(+0.29, –0.24)	46	2.68	–
J1845–0743	18:45:57.1763(9)	–07:43:38.16(5)	9.55157970586(+5, –4)	–33.444(±1)	(+1.5, –2.4)	72	2.46	–
J1847–0402	18:47:22.850(1)	–04:02:14.70(7)	1.67277577163(+6, –4)	–144.6391(+2, –3)	(+0.36, –0.0)	127	3.53	–
J1848–0123	18:48:23.596(1)	–01:23:58.48(6)	1.51644857592(+3, –6)	–11.9808(+4, –2)	0.19^{+2}_{-1}	111	3.49	–
J1849–0636	18:49:06.4647(2)	–06:37:06.91(1)	0.689011346415(±3)	–21.9601(±1)	(+0.02, –0.23)	108	3.23	–
J1852–0635	18:52:57.448(5)	–06:36:00.45(2)	1.90782618111(+5, –4)	–53.2438(±1)	(+0.59, –0.23)	109	3.25	–
J1852–2610	18:52:59.471(5)	–26:10:13.6(6)	2.9732067598(±6)	–0.7704(+1, –2)	(+1.5, –1.9)	25	2.43	–
J1857+0212	18:57:43.654(8)	+02:12:41.0(3)	2.40470716439(+9, –1)	–232.7342(±3)	(+1.4, –0.35)	71	2.72	–
J1900–2600	19:00:47.542(5)	–26:00:44.8(6)	1.63342812459(±1)	–0.54862(+4, –5)	(+0.15, –0.167)	40	3.44	–
J1901+0331	19:01:31.76(1)	+03:31:06.73(4)	1.52565744968(±2)	–17.3341(+8, –6)	(+2.3, –2.6)	127	3.13	–
J1901+0716	19:01:39.02(1)	+07:16:33.6(5)	1.55279458439(±4)	–5.577(±1)	(+4.6, –1.3)	28	2.61	–
J1901–0906	19:01:53.007(3)	–09:06:10.9(2)	0.561189668479(+6, –7)	–0.516(±2)	(+0.08, –0.19)	52	3.12	–
J1902+0556	19:02:42.60(1)	+05:56:25.8(1)	1.33943019014(±1)	–23.0839(±4)	(+1.3, –2.6)	32	2.44	–
J1902+0615	19:02:50.277(3)	+06:16:33.41(7)	1.48476895959(±4)	–16.9975(±1)	(+0.5, –2.3)	45	2.62	–
J1903+0135	19:03:29.973(1)	+01:35:38.73(4)	1.37116475483(±2)	–7.57312(±8)	(+0.16, –0.28)	94	3.18	–
J1903–0632	19:03:37.934(2)	–06:32:21.52(9)	2.31540809127(+7, –8)	–18.1244(+4, –3)	(+0.16, –0.71)	69	3.14	–
J1905–0056	19:05:27.752(6)	–00:56:40.8(3)	1.55476661905(+7, –6)	–7.3951(±1)	(+2.0, –1.7)	29	2.45	–
J1909+0007	19:09:35.252(2)	+00:07:56.84(9)	0.983329997648(+4, –7)	–5.3391(+3, –2)	(+0.37, –0.68)	71	3.07	–
J1909+0254	19:09:38.311(2)	+02:54:50.36(9)	1.01026940483(±1)	–5.61185(±4)	(+0.24, –0.1)	51	3.19	–
J1909+1102	19:09:48.6829(9)	+11:02:03.044(3)	3.5255695764(+1, –9)	–32.8178(±4)	$1^{+0.7}_{-1.1}$	95	3.19	–
J1909–3744	19:09:47.42783(7)	–37:44:14.767(3)	339.315686856(±5)	–1.6153(±2)	(+4, –5)	68	3.54	PPTA, B
J1910–0309	19:10:29.692(2)	–03:09:54.1(1)	1.98174395507(+3, –4)	–8.61183(±1)	-0.22^{+2}_{-4}	43	3.09	–
J1910+0358	19:10:09.016(3)	+03:58:30(1)	0.429135601854(+1, –8)	–0.8134(±3)	(+1.6, –1.9)	44	2.39	–
J1913–0440	19:13:54.1624(9)	–04:40:47.56(4)	1.21074218559(±2)	–5.9681(±1)	(+0.38, –0.21)	88	3.52	–
J1913+1400	19:13:24.352(1)	+14:00:52.50(3)	1.91764388163(+2, –3)	–2.95953(±7)	(+0.29, –0.24)	66	2.63	–

Table A.1 (Continued)

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	ν (Hz)	$\dot{\nu}$ (10^{-15} s^{-2})	$\ddot{\nu}$ (10^{-24} s^{-3})	N_{ToA}	T (yr)	Flags
J1915+1009	19:15:29.993(1)	+10:09:43.58(3)	2.47187153893(± 3)	-93.22115(+9, -8)	(+0.71, -0.2)	105	3.13	-
J1916+0951	19:16:32.333(1)	+09:51:25.97(3)	3.70019376623(± 8)	-34.52(+4, -3)	(+0.1, -0.51)	86	3.17	-
J1916+1312	19:16:58.67(2)	+13:12:50.0(4)	3.548050079(± 2)	-46.069(+4, -5)	(+49.9, -9.6)	39	161.67	-
J1917+1353	19:17:39.794(3)	+13:53:57.16(8)	5.13779943103(± 4)	-189.936(± 1)	$1^{+1.1}_{-1.5}$	71	2.47	-
J1919+0021	19:19:50.670(2)	+00:21:39.8(1)	0.78599927(+8, -7)	-4.74138(+2, -3)	(+0.18, -0.15)	97	3.26	-
J1926+0431	19:26:24.472(2)	+04:31:31.54(8)	0.931029279866(+1, -2)	-2.13409(+4, -5)	(+0.18, -0.1)	130	3.18	-
J1932+1059	19:32:14.038(2)	+10:59:33.21(5)	4.41464565156(± 1)	-22.5369(+7, -6)	(+1.6, -1.4)	116	3.64	-
J1932-3655	19:32:06.1280(6)	-36:55:01.78(3)	1.75002463079(+9, -1)	-0.8767(± 3)	(+0.6, -3.6)	39	2.38	-
J1935+1616	19:35:47.8255(2)	+16:16:39.723(4)	2.78750145981(+3, -5)	-46.6373(+1, -8)	(+0.51, -0.31)	59	2.15	-
J1941-2602	19:41:00.4169(1)	-26:02:05.884(9)	2.48226091399(+9, -8)	-5.89424(± 2)	(+0.21, -0.098)	105	3.09	-
J1943-1237	19:43:25.461(3)	-12:37:42.9(2)	1.02835150982(+1, -9)	-1.75624(+4, -3)	(+1.4, -0.2)	54	3.22	-
J1945-0040	19:45:28.33(3)	-00:40:59(1)	0.9563585837(± 1)	-0.47(± 2)	(+79.1, -88.5)	64	1.3	-
J1946-2913	19:46:51.757(5)	-29:13:48.1(3)	1.04226478935(± 1)	-1.61748(+5, -4)	(+0.2, -0.14)	63	3.27	-
J2006-0807	20:06:16.365(4)	-08:07:02.16(2)	1.72155151633(± 5)	-0.1355(+1, -9)	(+0.52, -0.4)	255	3.42	-
J2033+0042	20:33:31.12(2)	+00:42:24.1(9)	0.199465428208(+3, -2)	-0.38564(+6, -7)	(+0.21, -0.09)	129	3.28	-
J2038-3816	20:38:54.36(3)	-38:16:15.6(9)	0.633999188596(± 9)	-1.6728(± 3)	(+2.6, -6.7)	48	2.4	-
J2046-0421	20:46:00.1760(2)	-04:21:26.3(1)	0.646437789195(+1, -9)	-0.61473(+3, -2)	(+0.044, -0.064)	141	3.41	-
J2046+1540	20:46:39.336(5)	+15:40:33.6(1)	0.878513972444(± 3)	-0.14056(± 7)	(+0.21, -0.24)	72	3.39	-
J2048-1616	20:48:35.74(2)	-16:16:45(1)	0.509792367545(± 6)	-2.84929(± 2)	(+0.071, -0.055)	105	4.04	-
J2051-0827	20:51:07.52058(5)	-08:27:37.61(2)	221.796283548(± 3)	-0.6248(+9, -7)	(+2.6, -0.5)	193	3.23	B
J2053-7200	20:53:47.280(4)	-72:00:42.48(2)	2.9296611297(± 3)	-1.69606(+9, -8)	(+0.1, -0.19)	64	3.22	-
J2116+1414	21:16:13.761(1)	+14:14:20.38(4)	2.27193569866(± 4)	-1.49439(± 8)	(+0.13, -0.38)	127	3.18	-
J2129-5721	21:29:22.77664(9)	-57:21:14.2954(7)	268.359226956(± 2)	-1.5024(± 5)	(+14.1, -15.3)	100	2.22	PPTA, B
J2144-3933	21:44:12.01(1)	-39:33:58.4(3)	0.117511188481(+4, -5)	-0.0064(+1, -2)	(+0.038, -0.089)	95	3.2	-
J2145-0750	21:45:50.4552(8)	-07:50:18.56(3)	62.2958878113(± 3)	-0.111(± 1)	(+2.5, -0.1)	162	3.12	PPTA, B
J2155-3118	21:55:13.64(1)	-31:18:53.8(2)	0.97087088287(+3, -2)	-1.16876(+7, -8)	(+0.25, -0.21)	61	3.08	-
J2222-0137	22:22:05.96713(1)	-01:37:15.731(5)	30.4712133291(± 1)	-0.0535(+3, -5)	(+0.67, -0.17)	216	3.2	B
J2241-5236	22:41:42.03154(6)	-52:36:36.2491(6)	457.310149438(+9, -1)	-1.4408(± 6)	(+0.65, -0.81)	295	3.27	PPTA, B
J2248-0101	22:48:26.884(6)	-01:01:48.0(2)	2.09541027394(± 5)	-2.8961(± 1)	(+0.8, -2.5)	172	2.54	-
J2324-6054	23:24:27.14(1)	-60:54:05.794(9)	0.425987202198(+1, -9)	-0.46843(+3, -4)	(+0.061, -0.12)	87	3.07	-
J2330-2005	23:30:26.986(2)	-20:05:29.75(7)	0.608411174931(± 4)	-1.71419(± 1)	(+0.042, -0.097)	172	3.52	-
J2346-0609	23:46:50.54(1)	-06:10:01.04(4)	0.846407381972(± 3)	-0.9728(± 1)	(+0.04, -0.46)	236	3.23	-

Table A.2 List of the preferred timing noise model, Bayes factor when compared to the next simplest model and associated red noise parameters (errors indicate the 95 percent confidence intervals) for each pulsar in our data set. MSPs are indicated by a \star and clock reference pulsars by a \dagger . The red noise of clock reference pulsars (such as PSR J0437–4715) are contaminated by residual clock jumps, so should be used with caution. The full table can be found in the supplementary material.

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J0030+0451 \star	WTN	–	–	–
J0134–2937	WTN	–	–	–
J0151–0635	WTN	–	–	–
J0152–1637	WTN	–	–	–
J0206–4028	WTN	–	–	–
J0255–5304	WTN	–	–	–
J0348+0432 \star	WTN	–	–	–
J0401–7608	PLRN	10.6	$-10.2^{+0.2}_{-0.5}$	$4.0^{+14.4}_{-2.5}$
J0418–4154	WTN	–	–	–
J0437–4715 $\star\dagger$	PLRN	4.2	$-10.8^{+0.7}_{-0.4}$	$9.5^{+3.8}_{-2.5}$
J0450–1248	WTN	–	–	–
J0452–1759	PLRN	22.0	$-10.4^{+0.3}_{-0.7}$	$3.1^{+3.5}_{-2.0}$
J0525+1115	WTN	–	–	–
J0529–6652	WTN	–	–	–
J0533+0402	WTN	–	–	–
J0536–7543	WTN	–	–	–
J0601–0527	WTN	–	–	–
J0624–0424	WTN	–	–	–
J0627+0706	PLRN	60.2	$-10.0^{+0.4}_{-0.4}$	$3.5^{+2.5}_{-1.7}$
J0630–2834	WTN	–	–	–
J0646+0905	WTN	–	–	–
J0659+1414	PLRN	26.8	$-10.1^{+0.4}_{-3.5}$	$5.6^{+14.3}_{-2.1}$
J0711–6830 \star	WTN	–	–	–
J0729–1836	PLRN	191.6	$-9.7^{+0.3}_{-0.3}$	$6.1^{+3.0}_{-2.0}$
J0737–3039A \star	WTN	–	–	–
J0738–4042	PLRN+F2	5.4	$-9.8^{+0.3}_{-0.2}$	$6.5^{+1.8}_{-1.4}$
J0742–2822	PLRN	512.3	$-9.0^{+0.2}_{-0.1}$	$4.8^{+1.4}_{-1.0}$
J0758–1528	PLRN	3.2	$-10.7^{+0.6}_{-3.9}$	$4.1^{+15.8}_{-1.8}$
J0809–4753	PLRN	3.5	$-11.1^{+1.2}_{-3.6}$	$5.7^{+14.2}_{-3.5}$
J0820–1350	WTN	–	–	–
J0820–4114	WTN	–	–	–
J0835–4510	PLRN	3173.2	$-8.2^{+0.2}_{-0.2}$	$8.6^{+0.9}_{-0.9}$
J0837+0610	WTN	–	–	–
J0837–4135	PLRN	138.6	$-11.7^{+0.8}_{-1.0}$	$7.5^{+5.5}_{-2.9}$
J0840–5332	WTN	–	–	–
J0842–4851	WTN	–	–	–
J0846–3533	WTN	–	–	–

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J0855–3331	WTN	–	–	–
J0856–6137	WTN	–	–	–
J0904–4246	WTN	–	–	–
J0904–7459	WTN	–	–	–
J0907–5157	PLRN	92.7	$-10.8^{+0.7}_{-3.4}$	$5.8^{+14.2}_{-2.1}$
J0908–1739	WTN	–	–	–
J0908–4913	PLRN	523.7	$-9.6^{+0.2}_{-0.2}$	$5.0^{+1.0}_{-0.8}$
J0909–7212	WTN	–	–	–
J0922+0638	PLRN	101.1	$-9.4^{+0.2}_{-0.2}$	$5.6^{+1.8}_{-2.0}$
J0924–5302	PLRN	279.9	$-9.3^{+0.3}_{-0.2}$	$4.5^{+1.3}_{-1.1}$
J0924–5814	WTN	–	–	–
J0934–5249	WTN	–	–	–
J0942–5552	PLRN	523.1	$-9.0^{+0.2}_{-0.2}$	$5.9^{+1.6}_{-1.1}$
J0942–5657	PLRN	26.5	$-13.1^{+2.6}_{-1.0}$	$17.4^{+2.6}_{-11.6}$
J0944–1354	WTN	–	–	–
J0953+0755	WTN	–	–	–
J0955–5304	WTN	–	–	–
J0959–4809	WTN	–	–	–
J1001–5507	PLRN	492.9	$-9.1^{+0.2}_{-0.1}$	$4.6^{+1.2}_{-0.8}$
J1003–4747	WTN	–	–	–
J1012–5857	WTN	–	–	–
J1013–5934	WTN	–	–	–
J1016–5345	WTN	–	–	–
J1017–5621	WTN	–	–	–
J1017–7156*	WTN	–	–	–
J1022+1001*	WTN	–	–	–
J1032–5911	WTN	–	–	–
J1034–3224	WTN	–	–	–
J1036–4926	WTN	–	–	–
J1041–1942	WTN	–	–	–
J1042–5521	WTN	–	–	–
J1043–6116	WTN	–	–	–
J1045–4509*	WTN	–	–	–
J1046–5813	PLRN	7.0	$-13.2^{+2.9}_{-1.1}$	$18.3^{+1.6}_{-14.4}$
J1047–6709	WTN	–	–	–
J1048–5832	PLRN	1258.2	$-8.6^{+0.2}_{-0.1}$	$6.3^{+1.2}_{-1.0}$
J1056–6258	PLRN	297.3	$-9.7^{+0.2}_{-0.2}$	$2.9^{+1.1}_{-1.0}$
J1057–5226	PLRN	267.5	$-9.9^{+0.2}_{-0.2}$	$5.9^{+2.3}_{-1.5}$
J1057–7914	WTN	–	–	–
J1059–5742	WTN	–	–	–
J1105–6107	PLRN	347.5	$-8.9^{+0.3}_{-0.2}$	$4.1^{+1.7}_{-1.2}$

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J1110–5637	PLRN	49.0	$-10.3^{+0.5}_{-0.9}$	$6.5^{+5.7}_{-3.1}$
J1112–6613	PLRN	35.2	$-9.6^{+0.3}_{-0.2}$	$3.5^{+2.8}_{-1.7}$
J1112–6926	WTN	–	–	–
J1114–6100	WTN	–	–	–
J1116–4122	PLRN	4.8	$-13.9^{+3.8}_{-0.8}$	$17.0^{+2.1}_{-5.2}$
J1121–5444	PLRN	111.1	$-9.8^{+0.3}_{-0.3}$	$6.0^{+3.6}_{-2.1}$
J1123–6259	WTN	–	–	–
J1126–6942	WTN	–	–	–
J1133–6250	WTN	–	–	–
J1136+1551	PLRN	6.5	$-10.4^{+0.5}_{-3.9}$	$4.0^{+15.9}_{-2.1}$
J1136–5525	PLRN	174.3	$-10.2^{+0.5}_{-0.6}$	$7.7^{+4.3}_{-2.5}$
J1141–3322	WTN	–	–	–
J1141–6545	PLRN	186.8	$-10.3^{+0.4}_{-0.5}$	$4.7^{+3.2}_{-1.9}$
J1146–6030	WTN	–	–	–
J1157–6224	PLRN	97.5	$-10.1^{+0.2}_{-0.2}$	$3.3^{+1.8}_{-1.4}$
J1202–5820	PLRN	69.2	$-10.3^{+0.4}_{-0.6}$	$5.3^{+4.0}_{-2.0}$
J1210–5559	PLRN	4.5	$-14.7^{+3.2}_{-1.4}$	$13.8^{+6.1}_{-10.4}$
J1224–6407	PLRN	372.6	$-10.7^{+0.2}_{-0.2}$	$5.0^{+1.9}_{-1.5}$
J1231–6303	WTN	–	–	–
J1239–6832	WTN	–	–	–
J1243–6423	PLRN	950.4	$-10.2^{+0.2}_{-0.2}$	$4.5^{+1.0}_{-0.8}$
J1253–5820	PLRN	67.0	$-10.7^{+0.5}_{-0.7}$	$5.6^{+4.5}_{-2.7}$
J1257–1027	WTN	–	–	–
J1259–6741	WTN	–	–	–
J1305–6455	PLRN	197.0	$-9.7^{+0.3}_{-0.6}$	$4.9^{+3.0}_{-1.6}$
J1306–6617	PLRN	7.8	$-11.5^{+1.6}_{-2.6}$	$16.5^{+3.4}_{-12.3}$
J1312–5402	WTN	–	–	–
J1312–5516	WTN	–	–	–
J1319–6056	PLRN	24.8	$-10.3^{+0.4}_{-3.7}$	$3.1^{+15.5}_{-3.4}$
J1320–5359	PLRN	43.1	$-13.8^{+3.1}_{-0.9}$	$19.4^{+0.6}_{-13.4}$
J1326–5859	PLRN	718.2	$-10.1^{+0.3}_{-0.2}$	$5.4^{+1.3}_{-1.0}$
J1326–6408	WTN	–	–	–
J1326–6700	PLRN	107.9	$-9.3^{+0.2}_{-0.2}$	$3.5^{+1.8}_{-1.4}$
J1327–6222	PLRN	946.9	$-9.1^{+0.2}_{-0.2}$	$4.2^{+1.1}_{-1.0}$
J1327–6301	WTN	–	–	–
J1328–4357	PLRN	11.0	$-13.2^{+3.0}_{-0.8}$	$19.9^{+0.0}_{-15.4}$
J1338–6204	WTN	–	–	–
J1350–5115	WTN	–	–	–
J1355–5153	PLRN	4.1	$-13.4^{+2.9}_{-1.0}$	$16.4^{+3.5}_{-11.8}$
J1356–5521	WTN	–	–	–
J1359–6038	PLRN	1556.6	$-10.0^{+0.2}_{-0.1}$	$5.1^{+0.9}_{-0.8}$

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J1401–6357	PLRN	693.7	$-9.8^{+0.3}_{-0.3}$	$7.5^{+2.9}_{-2.1}$
J1413–6307	PLRN	143.4	$-9.4^{+0.4}_{-0.3}$	$4.6^{+3.0}_{-2.4}$
J1418–3921	WTN	–	–	–
J1420–5416	WTN	–	–	–
J1424–5822	WTN	–	–	–
J1428–5530	WTN	–	–	–
J1430–6623	PLRN	26.4	$-11.0^{+0.2}_{-0.0}$	$3.2^{+7.2}_{-0.2}$
J1435–5954	WTN	–	–	–
J1452–6036	WTN	–	–	–
J1453–6413	PLRN	156.9	$-10.9^{+0.2}_{-0.2}$	$3.2^{+1.5}_{-1.3}$
J1456–6843	WTN	–	–	–
J1457–5122	WTN	–	–	–
J1507–4352	PLRN	16.2	$-10.3^{+0.4}_{-0.9}$	$3.6^{+5.5}_{-2.5}$
J1507–6640	WTN	–	–	–
J1511–5414	WTN	–	–	–
J1512–5759	PLRN	254.7	$-9.9^{+0.3}_{-0.4}$	$6.8^{+3.9}_{-2.4}$
J1514–4834	WTN	–	–	–
J1522–5829	PLRN	28.9	$-12.3^{+2.1}_{-1.6}$	$11.0^{+9.0}_{-6.3}$
J1527–3931	WTN	–	–	–
J1527–5552	PLRN	11.8	$-10.0^{+0.2}_{-4.1}$	$17.1^{+2.8}_{-15.8}$
J1528–3146*	WTN	–	–	–
J1534–5334	WTN	–	–	–
J1534–5405	PLRN	37.8	$-9.7^{+0.3}_{-0.3}$	$6.1^{+12.2}_{-1.8}$
J1539–5626	PLRN	17.2	$-9.8^{+0.3}_{-0.2}$	$3.1^{+2.7}_{-2.4}$
J1542–5034	PLRN	10.0	$-11.5^{+2.4}_{-0.6}$	$15.3^{+4.6}_{-11.1}$
J1543+0929	WTN	–	–	–
J1544–5308	WTN	–	–	–
J1549–4848	WTN	–	–	–
J1553–5456	WTN	–	–	–
J1555–3134	WTN	–	–	–
J1557–4258	PLRN	10.7	$-12.3^{+1.5}_{-1.2}$	$8.7^{+5.0}_{-5.2}$
J1559–4438	PLRN	5.7	$-10.8^{+0.7}_{-2.7}$	$3.9^{+9.5}_{-0.7}$
J1559–5545	PLRN	221.5	$-8.8^{+0.2}_{-0.2}$	$4.9^{+1.6}_{-1.5}$
J1600–3053*	WTN	–	–	–
J1600–5044	PLRN	270.2	$-10.2^{+0.3}_{-0.4}$	$6.1^{+3.4}_{-2.1}$
J1603–2531	WTN	–	–	–
J1603–2712	WTN	–	–	–
J1603–7202*	WTN	–	–	–
J1604–4909	PLRN	133.1	$-10.3^{+0.3}_{-0.5}$	$5.4^{+2.9}_{-1.4}$
J1605–5257	WTN	–	–	–
J1613–4714	WTN	–	–	–

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J1622–4950	PLRN	211.8	$-4.9^{+0.6}_{-0.4}$	$7.3^{+3.4}_{-3.6}$
J1623–0908	WTN	–	–	–
J1623–4256	PLRN	21.2	$-12.6^{+2.3}_{-1.6}$	$19.9^{+0.0}_{-14.2}$
J1626–4537	WTN	–	–	–
J1633–4453	WTN	–	–	–
J1633–5015	WTN	–	–	–
J1639–4604	WTN	–	–	–
J1644–4559	PLRN	2519.7	$-9.9^{+0.2}_{-0.2}$	$6.2^{+1.0}_{-0.9}$
J1646–6831	WTN	–	–	–
J1651–4246	PLRN	125.1	$-12.8^{+2.5}_{-1.2}$	$19.7^{+0.2}_{-12.9}$
J1651–5222	WTN	–	–	–
J1651–5255	PLRN	23.6	$-12.2^{+2.5}_{-1.3}$	$12.6^{+7.3}_{-8.1}$
J1652–2404	WTN	–	–	–
J1700–3312	WTN	–	–	–
J1701–3726	WTN	–	–	–
J1703–1846	WTN	–	–	–
J1703–3241	WTN	–	–	–
J1703–4851	WTN	–	–	–
J1705–1906	PLRN	58.2	$-10.5^{+0.4}_{-0.6}$	$4.8^{+3.9}_{-2.1}$
J1705–3423	PLRN	27.3	$-10.8^{+0.6}_{-3.6}$	$6.2^{+13.7}_{-2.7}$
J1707–4053	WTN	–	–	–
J1708–3426	WTN	–	–	–
J1709–1640	PLRN	48.5	$-9.8^{+0.2}_{-0.4}$	$3.9^{+14.1}_{-1.4}$
J1709–4429	PLRN	504.8	$-9.1^{+0.3}_{-0.4}$	$8.0^{+2.6}_{-1.6}$
J1711–5350	PLRN	8.4	$-12.6^{+2.9}_{-1.0}$	$14.7^{+5.3}_{-11.0}$
J1715–4034	WTN	–	–	–
J1717–3425	PLRN	10.9	$-9.4^{+0.3}_{-2.9}$	$6.2^{+13.7}_{-3.3}$
J1717–4054	PLRN	40.7	$-11.0^{+1.6}_{-2.9}$	$11.4^{+8.6}_{-6.6}$
J1720–1633	WTN	–	–	–
J1720–2933	WTN	–	–	–
J1722–3207	PLRN	6.6	$-13.8^{+3.1}_{-0.8}$	$18.0^{+2.0}_{-13.3}$
J1722–3712	PLRN	372.6	$-9.0^{+0.2}_{-0.2}$	$4.2^{+1.1}_{-0.9}$
J1727–2739	WTN	–	–	–
J1730–2304*	WTN	–	–	–
J1731–4744	PLRN	195.4	$-9.5^{+0.2}_{-0.2}$	$5.0^{+1.5}_{-1.4}$
J1733–2228	WTN	–	–	–
J1736–2457	WTN	–	–	–
J1739–2903	PLRN	19.7	$-13.1^{+2.7}_{-1.1}$	$19.5^{+0.4}_{-14.5}$
J1740–3015	PLRN	128.8	$-8.9^{+0.3}_{-0.2}$	$5.2^{+1.0}_{-1.4}$
J1741–3927	PLRN	159.5	$-9.8^{+0.3}_{-0.3}$	$6.3^{+2.5}_{-1.6}$
J1743–3150	WTN	–	–	–

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J1745–3040	PLRN	68.7	$-14.3^{+3.3}_{-0.7}$	$18.9^{+0.5}_{-5.2}$
J1751–4657	WTN	–	–	–
J1752–2806	PLRN	292.9	$-9.6^{+0.2}_{-0.1}$	$2.9^{+0.8}_{-0.7}$
J1757–2421	WTN	–	–	–
J1759–2205	PLRN	54.9	$-10.1^{+0.3}_{-0.3}$	$4.4^{+2.0}_{-1.6}$
J1759–3107	WTN	–	–	–
J1801–0357	WTN	–	–	–
J1801–2920	WTN	–	–	–
J1803–2137	PLRN	41.3	$-8.6^{+0.5}_{-0.5}$	$17.9^{+2.0}_{-9.8}$
J1805–1504	WTN	–	–	–
J1807–0847	WTN	–	–	–
J1807–2715	PLRN	5.2	$-12.2^{+2.1}_{-1.9}$	$17.0^{+3.0}_{-12.8}$
J1808–0813	WTN	–	–	–
J1809–2109	WTN	–	–	–
J1810–5338	WTN	–	–	–
J1816–2650	WTN	–	–	–
J1818–1422	WTN	–	–	–
J1820–0427	PLRN	99.8	$-10.1^{+0.3}_{-0.5}$	$5.7^{+3.3}_{-2.0}$
J1822–2256	WTN	–	–	–
J1823–0154	WTN	–	–	–
J1823–1115	WTN	–	–	–
J1823–3106	PLRN	33.1	$-10.3^{+0.3}_{-3.5}$	$4.3^{+14.6}_{-4.9}$
J1824–0127	WTN	–	–	–
J1824–1945	PLRN	327.7	$-10.0^{+0.2}_{-0.2}$	$6.1^{+1.2}_{-1.2}$
J1825–0935	PLRN	570.6	$-9.0^{+0.2}_{-0.2}$	$4.9^{+2.1}_{-1.0}$
J1827–0750	PLRN	30.5	$-10.9^{+1.4}_{-1.4}$	$18.8^{+1.1}_{-13.2}$
J1829–1751	PLRN	187.9	$-9.6^{+0.2}_{-0.2}$	$6.0^{+2.6}_{-1.7}$
J1830–1135	WTN	–	–	–
J1832–0827	PLRN	67.8	$-10.4^{+0.6}_{-1.4}$	$6.1^{+7.9}_{-3.0}$
J1833–0338	PLRN	254.7	$-9.6^{+0.2}_{-0.2}$	$5.3^{+1.6}_{-1.1}$
J1833–0827	PLRN	28.4	$-11.9^{+1.8}_{-1.7}$	$15.4^{+4.5}_{-10.5}$
J1834–0426	WTN	–	–	–
J1835–1020	PLRN	84.7	$-8.9^{+0.3}_{-0.1}$	$3.5^{+3.2}_{-1.3}$
J1836–0436	WTN	–	–	–
J1836–1008	PLRN	85.4	$-10.3^{+0.7}_{-0.9}$	$7.8^{+3.5}_{-4.1}$
J1837–0653	WTN	–	–	–
J1840–0809	WTN	–	–	–
J1840–0815	WTN	–	–	–
J1841–0425	WTN	–	–	–
J1841+0912	PLRN	19.3	$-9.8^{+0.4}_{-0.6}$	$4.3^{+4.2}_{-2.5}$
J1842–0359	WTN	–	–	–

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J1843-0000	PLRN	6.3	$-12.7^{+3.0}_{-1.0}$	$19.0^{+0.9}_{-15.2}$
J1844-0433	WTN	—	—	—
J1845-0743	WTN	—	—	—
J1847-0402	PLRN	15.1	$-10.5^{+0.5}_{-4.3}$	$4.3^{+15.6}_{-1.8}$
J1848-0123	PLRN	46.7	$-10.7^{+0.7}_{-3.3}$	$5.7^{+12.0}_{-1.5}$
J1849-0636	PLRN	13.7	$-12.7^{+2.6}_{-1.5}$	$18.2^{+1.7}_{-14.3}$
J1852-0635	WTN	—	—	—
J1852-2610	WTN	—	—	—
J1857+0212	WTN	—	—	—
J1900-2600	WTN	—	—	—
J1901+0331	PLRN	277.3	$-9.5^{+0.2}_{-0.2}$	$4.4^{+1.3}_{-1.3}$
J1901+0716	PLRN	4.1	$-11.4^{+2.3}_{-1.3}$	$17.8^{+2.2}_{-13.9}$
J1901-0906	WTN	—	—	—
J1902+0556	WTN	—	—	—
J1902+0615	WTN	—	—	—
J1903+0135	PLRN	74.4	$-13.4^{+2.5}_{-0.8}$	$20.0^{+0.0}_{-12.4}$
J1903-0632	PLRN	10.2	$-10.1^{+0.3}_{-4.0}$	$2.7^{+16.4}_{-8.1}$
J1905-0056	WTN	—	—	—
J1909+0007	PLRN	64.3	$-10.3^{+0.3}_{-0.5}$	$6.6^{+10.9}_{-1.6}$
J1909+0254	WTN	—	—	—
J1909+1102	PLRN	183.9	$-10.6^{+0.5}_{-0.5}$	$7.9^{+3.5}_{-2.5}$
J1909-3744 [†]	WTN	—	—	—
J1910-0309	WTN	—	—	—
J1910+0358	WTN	—	—	—
J1913-0440	PLRN	175.6	$-10.9^{+0.6}_{-0.7}$	$7.5^{+5.0}_{-2.4}$
J1913+1400	WTN	—	—	—
J1915+1009	WTN	—	—	—
J1916+0951	PLRN	19.7	$-13.2^{+2.7}_{-1.2}$	$16.5^{+3.5}_{-12.2}$
J1916+1312	PLRN	111.7	$-9.3^{+0.3}_{-0.3}$	$6.0^{+3.3}_{-1.8}$
J1917+1353	PLRN	74.0	$-10.1^{+0.3}_{-0.2}$	$3.7^{+1.9}_{-1.4}$
J1919+0021	WTN	—	—	—
J1926+0431	WTN	—	—	—
J1932+1059	PLRN	206.5	$-10.4^{+0.3}_{-0.2}$	$5.4^{+2.5}_{-1.6}$
J1932-3655	WTN	—	—	—
J1935+1616	PLRN	31.3	$-10.8^{+0.3}_{-2.4}$	$4.5^{+15.5}_{-1.4}$
J1941-2602	WTN	—	—	—
J1943-1237	WTN	—	—	—
J1945-0040	WTN	—	—	—
J1946-2913	WTN	—	—	—
J2006-0807	WTN	—	—	—
J2033+0042	WTN	—	—	—

Table A.2 (Continued)

PSRJ	Model	$\ln(\mathcal{B})$	$\log_{10}(A)$ ($\text{yr}^{3/2}$)	β
J2038–3816	WTN	–	–	–
J2046–0421	WTN	–	–	–
J2046+1540	WTN	–	–	–
J2048–1616	WTN	–	–	–
J2051–0827*	WTN	–	–	–
J2053–7200	WTN	–	–	–
J2116+1414	WTN	–	–	–
J2129–5721*	WTN	–	–	–
J2144–3933	WTN	–	–	–
J2145–0750*	PLRN	33.1	$-11.3^{+0.3}_{-0.3}$	$4.8^{+3.3}_{-2.7}$
J2155–3118	WTN	–	–	–
J2222–0137*	WTN	–	–	–
J2241–5236*	PLRN	8.7	$-12.1^{+0.3}_{-0.3}$	$0.4^{+2.0}_{-0.4}$
J2248–0101	WTN	–	–	–
J2324–6054	WTN	–	–	–
J2330–2005	WTN	–	–	–
J2346–0609	PLRN	49.1	$-12.7^{+2.4}_{-1.3}$	$19.1^{+0.9}_{-13.7}$

B

Hidden-Markov Model glitch detector

B.1 HMM recipe and parameters

A complete description of the HMM framework and its application to glitch detection is given in [Melatos et al. \(2020\)](#). Applying the HMM glitch detector involves choosing a number of parameters, many of which vary between data sets. Here we lay out the HMM parameter choices which have been made in the present work.

Three essential groups of parameters are needed: those which specify the allowed $(\nu, \dot{\nu})$ states, those which relate the observations (ToAs) back to the $(\nu, \dot{\nu})$ states, and those for the probabilities of transitions between $(\nu, \dot{\nu})$ states. The choices of parameters which specify how the ToAs can be related to the $(\nu, \dot{\nu})$ states follow [Melatos et al. \(2020\)](#); we refer the reader to Section 3.3 and Appendix C of that paper for further discussion.

Specifying the allowed $(\nu, \dot{\nu})$ parameters amounts to specifying a discretisation of a region of $(\nu, \dot{\nu})$ space. We first note that the hidden $(\nu, \dot{\nu})$ states should be thought of as a deviation away from a secular phase model, where this secular phase model is derived from an initial TEMPO2 fit. As such, the region of $(\nu, \dot{\nu})$ space to be discretised is typically taken to be a region containing $(0 \text{ Hz}, 0 \text{ Hz s}^{-1})$. As a default, we take the boundaries in the ν direction to be $[\nu_-, \nu_+] = [-3 \times 10^{-7}, 3 \times 10^{-7}] \text{ Hz}$, and the spacing between discrete states to be $\eta_\nu = 10^{-9} \text{ Hz}$. The default boundaries and spacing in the $\dot{\nu}$ direction are $[\dot{\nu}_-, \dot{\nu}_+] = [-10^{-15}, 10^{-15}] \text{ Hz s}^{-1}$ and $\eta_{\dot{\nu}} = 2 \times 10^{-16} \text{ Hz s}^{-1}$ respectively. In the course of the analyses presented here, we found it necessary in some cases and advantageous in others to modify the boundaries of the $\dot{\nu}$ region to better account for the timing noise present in each pulsar. When the timing noise is strong enough that $\dot{\nu}$ wanders outside the default region, we enlarge the $\dot{\nu}$ region to allow the HMM to track the evolution of $\dot{\nu}$. In contrast, for some pulsars the timing noise is small enough that the scale of wandering of $\dot{\nu}$ is only a small fraction of the default $\dot{\nu}$ region. In this case it is advantageous (but

not necessary) to reduce the size of the $\dot{\nu}$ region, which allows for faster computation and more stringent upper limits on the size of undetected glitches.

Finally, we consider the probabilities of transitions between $(\nu, \dot{\nu})$ states. There is a trivial element: over a gap between consecutive ToAs of length z , in the absence of noise the state $(\nu, \dot{\nu})$ should transition to $(\nu + \dot{\nu}z, \dot{\nu})$. However, timing noise is not negligible in these pulsars. In general timing noise is incorporated into the HMM through a model which assumes the presence of some kind of stochastic term in the phase model, which causes wandering in $(\nu, \dot{\nu})$. For most of the pulsars in this study we follow the prescription of Melatos et al. (2020), which assumes that there is a white noise term in the second time derivative of ν :

$$\frac{d^2\nu}{dt^2} = \xi(t), \quad (\text{B.1})$$

where $\xi(t)$ satisfies

$$\langle \xi(t) \rangle = 0, \quad (\text{B.2})$$

$$\langle \xi(t)\xi(t') \rangle = \sigma^2\delta(t-t'). \quad (\text{B.3})$$

From these equations we may calculate the covariance matrix Σ of ν and $\dot{\nu}$ over a ToA gap of length z :

$$\Sigma = \sigma^2 \begin{pmatrix} z^3/3 & z^2/2 \\ z^2/2 & z \end{pmatrix}. \quad (\text{B.4})$$

Once Σ is given, the probabilities of transitions between $(\nu, \dot{\nu})$ states are given by equations (10)–(12) of Melatos et al. (2020). It is important to recognise that (B.1) is not proposed as a physical model specific to an individual pulsar. Rather, equation (B.1) is simply a generic, mathematically precise means to introduce stochastic fluctuations into the phase model, that are qualitatively consistent with random-walk character of the phase residuals observed in pulsars as a class. It has been shown to perform reliably on synthetic and real data in previous studies (Melatos et al., 2020).

The σ parameter controls the strength of the timing noise which is included in the HMM, and currently we do not have a reliable method of making a measurement of the timing noise in each pulsar and converting that to a suitable value of σ . As such, for this work we adopt a rule of thumb in Melatos et al. (2020) which sets a minimum value of σ based on the spacing in the $\dot{\nu}$ grid $\eta_{\dot{\nu}}$:

$$\sigma = \eta_{\dot{\nu}} \langle z \rangle^{-1/2}. \quad (\text{B.5})$$

In a few cases we find that the above model is not satisfactory: the analysis produces an implausibly large number of glitch candidates, none of which show any signature in the timing residuals. In these cases, we adopt an alternative form for Σ , which is predicated on a white noise term in the *first* time derivative of ν :

$$\frac{d\nu}{dt} = \zeta(t), \quad (\text{B.6})$$

with $\zeta(t)$ satisfying equations analogous to (B.2) and (B.3). In this case Σ has the simpler form

$$\Sigma = \sigma^2 \begin{pmatrix} z & 0 \\ 0 & 1 \end{pmatrix} \quad (\text{B.7})$$

The choice of σ is no longer given by (B.5), but instead is set by hand to match the scale of variation in $\dot{\nu}$ which is observed in the $\dot{\nu}$ paths recovered by the HMM, which is typically on the order of $10^{-15} \text{ Hz s}^{-1}$.

B.2 Glitch upper-limits

We define $\Delta\nu_g^{90\%}$ in the following way: if a large number of synthetic data sets are produced, each with a glitch of size $\Delta\nu_g^{90\%}$ injected at a randomly chosen epoch, we expect that the HMM will detect a glitch in 90% of those cases at an epoch which is not more than one ToA away from the injected epoch.

This definition suggests an empirical method of calculating $\Delta\nu_g^{90\%}$:

1. Make an estimate for $\Delta\nu_g^{90\%}$, denoted $\Delta\nu_g^{x\%}$.
2. Generate a set of 100 synthetic data sets each with a glitch of size $\Delta\nu_g^{x\%}$.
3. Perform HMM analyses on each of the 100 synthetic data sets, and record the number n of data sets for which the HMM detects a glitch within one ToA of the injected epoch.
4. If $n = 90$, terminate and take $\Delta\nu_g^{90\%} = \Delta\nu_g^{x\%}$. Otherwise, choose an updated $\Delta\nu_g^{x\%}$ and return to step (ii).

The refinement of $\Delta\nu_g^{x\%}$ proceeds essentially as a binary search: we choose an initial possible range for $\Delta\nu_g^{90\%}$, typically $[\Delta\nu_g^-, \Delta\nu_g^+] = [10^{-9}, 10^{-7}] \text{ Hz}$ and take $\Delta\nu_g^{x\%}$ to bisect this range logarithmically (i.e. at the first iteration, $\Delta\nu_g^{x\%} = 10^{-8} \text{ Hz}$). Then, if the number of detected glitches exceeds 90, the range is refined to be $[\Delta\nu_g^-, \Delta\nu_g^+] = [\Delta\nu_g^-, \Delta\nu_g^{x\%}]$. Similarly if the number of detected glitches is less than 90, the range is

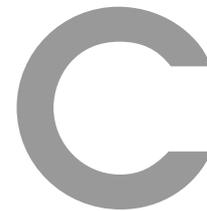
Table B.1 Pulsar averaged 90% upper-limits on $\Delta\nu_g/\nu$.

PSR	$\Delta\nu_g^{90\%}/\nu$ ($\times 10^{-9}$)								
J0614+2229	19	J0627+0706	2.7	J0631+1036	5.7	J0659+1414	1.1	J0729-1448	10
J0742-2822	4.6	J0835-4510	4.0	J0842-4851	5.1	J0855-4644	6.0	J0901-4624	4.3
J0908-4913	5.9	J0940-5428	2.4	J1003-4747	0.17	J1015-5719	4.1	J1016-5857	3.4
J1019-5749	1.6	J1028-5819	0.88	J1048-5832	2.0	J1052-5954	5.7	J1055-6028	3.5
J1057-5226	3.4	J1105-6107	6.3	J1112-6103	1.2	J1138-6207	5.6	J1248-6344	1.4
J1301-6305	13	J1320-5359	2.2	J1327-6400	8.2	J1341-6220	15	J1357-6429	3.8
J1359-6038	5.3	J1406-6121	3.1	J1410-6132	43	J1413-6141	6.8	J1420-6048	3.4
J1452-6036	0.34	J1524-5625	4.1	J1541-5535	14	J1600-5044	19	J1602-5100	30
J1614-5048	47	J1617-5055	4.0	J1626-4807	5.6	J1627-4706	13	J1644-4559	15
J1646-4346	12	J1650-4502	14	J1701-4533	0.57	J1702-4128	5.8	J1702-4310	5.3
J1705-3950	3.6	J1709-4429	1.3	J1716-4005	3.9	J1718-3825	0.55	J1721-3532	2.6
J1726-3530	77	J1730-3350	6.9	J1731-4744	5.1	J1734-3333	30	J1737-3137	4.5
J1740-3015	7.3	J1750-3157	2.2	J1757-2421	6.0	J1801-2304	3.1	J1803-2137	4.2
J1825-0935	9.9	J1826-1334	2.8	J1835-0643	0.23	J1837-0604	4.7	J1841-0345	2.0
J1841-0425	1.8	J1841-0524	4.8	J1844-0256	6.1	J1847-0402	1.9		

refined to $[\Delta\nu_g^-, \Delta\nu_g^+] = [\Delta\nu_g^{x\%}, \Delta\nu_g^+]$. A new choice of $\Delta\nu_g^{x\%}$ is then made to bisect the new range logarithmically, and the procedure repeats.

The generation of synthetic data sets in step (ii) is done using LIBSTEMPO¹. We inject additive Gaussian error at a level commensurate with the reported ToA error, but we do *not* inject red timing noise — injecting the latter tends to create difficulties in automating step (iii). While we do expect that timing noise in the data impacts the ability of the HMM to detect small glitches near the threshold of detectability, we do not expect that it makes a significant difference to the 90% frequentist upper limit. Results are given in Table B.1.

¹<https://github.com/vallis/libstempo>



Parkes timing results

Table C.1 Preferred timing models and associated Bayes factors compared to the standard PL model. For pulsars where the PL model is preferred, we list $\ln(\mathcal{B}_{\text{PL+F2}}^{\text{PL}})$

PSR	Preferred model	$\ln(\mathcal{B})$	PSR	Preferred model	$\ln(\mathcal{B})$	PSR	Preferred model	$\ln(\mathcal{B})$
J0614+2229	PL	2.9	J0627+0706	PL	5.8	J0631+1036	PL	-1.6
J0659+1414	PL+F2+PM	21.0	J0729-1448	PL	1.3	J0742-2822	PL	3.7
J0835-4510	PL+F2	227.2	J0842-4851	PL	6.0	J0855-4644	PL+F2	16.8
J0901-4624	PL+F2	122.2	J0908-4913	PL+F2+PM	11.8	J0940-5428	PL+F2	49.9
J1003-4747	PL+PM	15.3	J1015-5719	PL+F2	3.6	J1016-5857	PL+F2	49.7
J1019-5749	PL	8.2	J1028-5819	PL+F2	14.9	J1048-5832	PL+F2	65.0
J1052-5954	PL	-2.7	J1055-6028	PL	1.7	J1057-5226	PL+PM	3.6
J1105-6107	PL	1.3	J1112-6103	PL+F2	10.3	J1138-6207	PL	5.0
J1248-6344	PL	11.4	J1301-6305	PL+F2	88.0	J1320-5359	PL+F2+PM	14.2
J1327-6400	PL	4.4	J1341-6220	PL	5.3	J1357-6429	PL+F2	190.5
J1359-6038	PL+PM	0.8	J1406-6121	PL	5.8	J1410-6132	PL+F2	8.2
J1413-6141	PL	-0.5	J1420-6048	PL+F2+LFC	255.5	J1452-6036	PL	9.0
J1524-5625	PL+F2	73.6	J1541-5535	PL	-0.9	J1600-5044	PL	0.8
J1614-5048	PL+F2	56.1	J1617-5055	PL+F2	19.8	J1626-4807	PL	40.7
J1627-4706	PL	10.7	J1644-4559	PL	-1.2	J1646-4346	PL+F2	17.2
J1650-4502	PL	2.8	J1701-4533	PL	89.1	J1702-4128	PL+F2	8.8
J1702-4310	PL+F2	80.7	J1705-3950	PL	8.9	J1709-4429	PL+F2+PM	187.7
J1716-4005	PL	9.7	J1718-3825	PL+F2+PM	57.8	J1721-3532	PL	0.4
J1726-3530	PL+F2	30.5	J1730-3350	PL+F2	20.7	J1731-4744	PL+F2	5.4
J1734-3333	PL+F2	29.6	J1737-3137	PL+F2	5.5	J1740-3015	PL	3.0
J1750-3157	PL	65.2	J1757-2421	PL	5.5	J1801-2304	PL	2.9
J1801-2451	PL+F2	345.5	J1803-2137	PL+F2	102.6	J1825-0935	PL	7.9
J1826-1334	PL+F2+PM	33.2	J1835-0643	PL	13.1	J1837-0604	PL	-1.4
J1841-0345	PL	55.0	J1841-0425	PL+F2	14.5	J1841-0524	PL	7.4
J1844-0256	PL	9.2	J1847-0402	PL	5.4			

Table C.2 Inferred astrometric (RAJ, DECJ, μ_α , μ_δ) and rotational (ν , $\dot{\nu}$, $\ddot{\nu}$) parameters for all 51 glitching pulsars in our sample. All values are in reference to the MJD listed under PEPOCH. Uncertainties in parentheses indicate the 68 per cent confidence intervals scaled to the last significant figure. Asymmetric confidence intervals are individually listed. Lower- and upper-limits on $\dot{\nu}$, μ_α or μ_δ are given by two comma-separated values in parentheses. *Proper-motion fixed to value from [Dodson et al. \(2003\)](#).

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	PEPOCH (MJD)	ν (Hz)	$\dot{\nu}$ (10^{-14} s $^{-2}$)	$\ddot{\nu}$ (10^{-24} s $^{-3}$)	μ_α (mas yr $^{-1}$)	μ_δ (mas yr $^{-1}$)	N_{ToA}	T (yr)	MJD range (MJD)
J0631+1036	06:31:28(1)	10:37:11(4)	54750	3.47463459(3)	-126.249(9)	(-15, 21)	(-300, 700)	(-600, 600)	41	3.7	57165–58531
J0729–1448	07:29:16(1)	-14:48:40(2)	55297	3.97302972(3)	-177.76(4)	(1.1, 1.5)	(-320, 220)	(0, 700)	172	11.6	54220–58469
J0742–2822	07:42:49.0(3)	-28:22:43.8(4)	55352	5.99624469(2)	-60.44(2)	(-2.3, 8.8)	(-40, 110)	(-290, -130)	774	15.0	52988–58469
J0835–4510	08:35:20.6 $^{+0.8}_{-0.7}$	-45:10:33.5 $^{+0.5}_{-0.6}$	56364	11.1882333(4)	-1545.8(6)	950(5)	-49.68(6)*	29.9(1)*	414	11.5	54260–58469
J0901–4624	09:01:40.11(4)	-46:24:48.45(3)	55268	2.262304515(1)	-44.737(1)	1.19(9)	(-3, 7)	(-3, 9)	234	20.9	50849–58469
J0908–4913	09:08:35.5(3)	-49:13:06.4(2)	55332	9.36627103(2)	-132.56(1)	4.3(7)	-37(9)	31(10)	375	27.7	48860–58972
J0940–5428	09:40:58.3(2)	-54 : 28 : 40.2(1)	55335	11.42109977(7)	-427.41(7)	47.7 \pm 2.8	(-20, 12)	(-13, 19)	303	22.9	50849–59202
J1015–5719	10:15:38.0(3)	-57:19:12.1(2)	55332	7.147923371(5)	-292.619(7)	20(2)	(-80, 0)	(30, 130)	154	11.6	54220–58469
J1016–5857	10:16:21.3(5)	-58:57:11.3(2)	55369	9.3105496(3)	-695.5(2)	121(5)	(40, 120)	(0, 80)	352	19.6	51299–58469
J1019–5749	10:19:52.1 $^{+0.2}_{-0.1}$	-57:49:06.22(7)	55434	6.153628421(6)	-75.97(2)	(3.1, 4.5)	(-35, 17)	(-51, 5)	153	11.4	54302–58469
J1028–5819	10:28:27.9(1)	-58:19:06.21(6)	55459	10.94038305(1)	-192.82(2)	20(4)	(-47, -9)	(-36, 4)	137	10.7	54563–58469
J1048–5832	10:48:12.5(9)	-58:32:04.2(5)	55453	8.083519(2)	-617.9(5)	152(9)	(40, 140)	(70, 190)	523	28.9	47909–58469
J1052–5954	10:52:38.1(1)	-59:54:44.25 $^{+0.06}_{-0.07}$	55292	5.53716823(3)	-61.07(4)	(13, 19)	(-10, 70)	(-110, 30)	101	6.8	54220–56708
J1055–6028	10:55:39.3(6)	-60:28:35.5(3)	55397	10.033766397(7)	-297.08(3)	(6, 20)	(-60, 140)	(-30, 150)	171	10.9	54505–58469
J1105–6107	11:05:26.1(7)	-61:07:49.7(3)	55303	15.8230056(9)	-395.9(2)	(-9, 9)	(-100, 0)	(50, 150)	393	23.5	49868–58469
J1112–6103	11:12:14.8 $^{+0.7}_{-0.6}$	-61:03:30.9(3)	55456	15.390826(2)	-739.8(7)	149(21)	(-50, 70)	(-80, 40)	312	20.9	50849–58469
J1248–6344	12:48:46.4(2)	-63:44:09.37(7)	55392	5.0418245518(3)	-42.999(2)	(-2.1, 0.1)	(-120, -40)	(-110, -10)	90	6.8	54219–56709
J1301–6305	13:01:45.7(7)	-63:05:34.5(3)	55370	5.4164292(8)	-773.2(3)	278(7)	(0, 120)	(-50, 90)	274	20.5	50940–58444
J1320–5359	13:20:53.92(2)	-53:59:05.39(1)	55408	3.5747990403(6)	-11.812(1)	0.43(6)	13(2)	52(2)	282	21.7	50536–58469
J1341–6220	13:41:42 $^{+4}_{-5}$	-62:20:17 $^{+3}_{-2}$	56345	5.1690544(9)	-676.7(5)	(1.3, 3.1)	(-500, 500)	(-400, 600)	195	11.6	54220–58469
J1357–6429	13:57:02.5 $^{+0.9}_{-1.0}$	-64:29:30.2(5)	55000	6.0178443(6)	-1279.32(3)	1039(14)	(-100, 60)	(-210, -10)	293	19.1	51491–58469
J1406–6121	14:06:49.9 $^{+0.3}_{-0.2}$	-61:21:27.8 $^{+0.1}_{-0.2}$	55390	4.692757029(1)	-120.535(6)	(3.1, 7.5)	(-50, 110)	(-180, 20)	97	6.8	54220–56708
J1410–6132	14:10:22(1)	-61:32:00.5 $^{+0.9}_{-0.8}$	55433	19.9780635(5)	-1266.0(6)	180(26)	(-300, 140)	(-430, 130)	155	11.3	54353–58469
J1413–6141	14:13:10(2)	-61:41:15(1)	56011	3.4995041(6)	-408.2(2)	(-30, 56)	(-30, 390)	(-230, 250)	282	20.9	50849–58469
J1420–6048	14:20:08.2(5)	-60:48:17.5(3)	55404	14.661265(2)	-1748.7(9)	997(18)	(-20, 60)	(-80, 20)	344	19.5	51333–58469
J1452–6036	14:52:51.89(2)	-60:36:31.37(2)	55370	6.4519531394(5)	-6.0339(8)	(-0.1, 0.5)	-5(3)	-5(3)	150	11.6	54220–58469
J1524–5625	15:24:49.82 $^{+0.03}_{-0.04}$	-56:25:24.07(3)	55441	12.78268099(1)	-637.03(3)	137(2)	(-6, 12)	(-10, 18)	161	11.6	54220–58469
J1614–5048	16:14:11 $^{+6}_{-7}$	-50 : 48 : 02(7)	54359	4.313262(5)	-905(1)	272(5)	(-600, 200)	(-600, 400)	513	27.8	48329–58469
J1617–5055	16:17:29(5)	-50:55:11(5)	54450	14.4093408(3)	-2837.3(9)	2010(300)	(-800, 400)	(-600, 600)	137	6.8	54220–56708
J1644–4559	16:44:49.3(1)	-45:59:09.8(2)	57600	2.197422934(5)	-9.701(1)	(0.27, 0.43)	(-14, 6)	(-38, 14)	366	29.1	47913–58534

Table C.2 (Continued) *Position fixed to values from [Dexter et al. \(2017\)](#) where POSEPOCH is MJD 57259. †Position and proper-motion fixed to values from [Zeiger et al. \(2008\)](#) where POSEPOCH is MJD 53348.

PSRJ	RAJ (hh:mm:ss)	DECJ (°:':")	PEPOCH (MJD)	ν (Hz)	$\dot{\nu}$ (10^{-15} s^{-2})	$\ddot{\nu}$ (10^{-24} s^{-3})	μ_α (mas yr $^{-1}$)	μ_δ (mas yr $^{-1}$)	N_{ToA}	T (yr)	MJD range (MJD)
J1646–4346	16:46:50(1)	–43:45:53(2)	55388	4.3165103(1)	–208.43(8)	29(2)	(–150, 30)	(–230, 270)	380	28.9	47912–58469
J1650–4502	16:50:32.5 $^{+0.8}_{-0.9}$	–45:02:31(1)	55389	2.625528094(5)	–11.124(9)	(–1, 9)	(–390, –30)	(–800, 0)	138	11.6	54220–58469
J1702–4128	17:02:52.5(2)	–41:28:48.2(5)	55366	5.489885905(1)	–157.685(2)	5.8(3)	(–70, 30)	(–250, 50)	141	11.6	54220–58470
J1702–4310	17:02:26.94(6)	–43:10:41.5(1)7	55375	4.1563312(1)	–385.36(6)	49(2)	(–26, –10)	(–44, –2)	216	19.8	51222–58470
J1705–3950	17:05:29.8(4)	–39:50:58(1)	56462	3.1351208446(5)	–59.445(1)	(–0.45, –0.11)	(–300, –160)	(–250, 250)	158	13.7	53951–58972
J1709–4429	17:09:42.75(6)	–44:29:08.3(2)	55315	9.755973(1)	–869.3(3)	273(4)	17(4)	11(10)	441	28.9	47909–58470
J1718–3825	17:18:13.558(7)	–38:25:17.83(2)	55374	13.391416606(7)	–236.311(5)	21.3(2)	–11(1)	(–2, 6)	250	20.8	50877–58470
J1730–3350	17:30:32.4(5)	–33:50:34(2)	55400	7.1685489(5)	–433.8(1)	54(2)	(–200, –80)	(–820, –180)	274	21.3	50538–58321
J1731–4744	17:31:42.2(1)	–47:44:37.1(2)	55402	1.20497647(3)	–23.622(8)	2.5(1)	60(11)	–178(24)	317	25.8	49043–58470
J1734–3333	17:34:27(2)	–33:33:31 $^{+8}_{-7}$	55341	0.854866535(3)	–166.63(1)	3.9(6)	(–400, 400)	(–500, 700)	173	16.4	50686–56672
J1737–3137	17:37:04.3(2)	–31:37:26.4 $^{+1.0}_{-0.4}$	55292	2.21984993(1)	–68.30(2)	3.1(2)	(20, 100)	(190, 710)	166	13.0	54220–58972
J1740–3015	17:40:33.8 $^{+0.5}_{-0.4}$	–30:15:45(3)	55472	1.64767845(5)	–126.57(5)	(8, 20)	(–140, 120)	(180, 820)	195	11.6	54220–58470
J1757–2421	17:57:29.33(3)	–24:22:04(1)	55433	4.2715558534(8)	–23.570(1)	(0.24, 0.76)	–	–	155	12.3	53974–58470
J1801–2304	18:01:19.8149(6)*	–23:04:44.63(1)*	57259	2.40466187(2)	–652.95(9)	(2, 4)	(–33, 21)	(–825, 625)	157	12.3	53974–58470
J1801–2451	18:01:00.016(8)†	–24:51:27.5(2)†	55297	8.002151(1)	–7979(4)	304(7)	–11(9)*	–1(15)*	156	26.2	48896–58470
J1803–2137	18:03:51.4(2)	–21:37:03(5)	55430	7.48029082(3)	–748.63(9)	239(6)	(20, 100)	(–900, 500)	151	11.6	54220–58470
J1825–0935	18:25:30.6(4)	–09:35:21(2)	56550	1.30036246(9)	–8.88(3)	(–0.8, 2.2)	(–180, –40)	(–990, –490)	179	18.3	51844–58533
J1826–1334	18:26:13.19(6)	–13:34:46.5 $^{+0.4}_{-0.3}$	54286	9.85345601(3)	–730.50(3)	172(8)	32(9)	(–80, 40)	141	11.5	54220–58404
J1837–0604	18:37:43.4(6)	–06:04:49(2)	55558	10.38323993(8)	–485.9(1)	(–40, 240)	(–310, 230)	(–1000, 0)	106	7.5	53968–56708
J1841–0345	18:41:38.7 $^{+0.2}_{-0.1}$	–03:48:43.8 $^{+0.5}_{-0.6}$	54867	4.899957507(7)	–139.000(3)	(5, 15)	(–520, 20)	(–500, 300)	45	3.7	57165–58531
J1841–0524	18:41:49.3(1)	–05:24:30.4(3)	55362	2.24311056(1)	–117.55(1)	(–1, 9)	(–60, 80)	(–340, 200)	91	6.7	54268–56708
J1847–0402	18:47:22.84(5)	–04:02:14.6 $^{+0.1}_{-0.2}$	55448	1.6728026921(1)	–14.4652(2)	(0.08, 1.2)	(–6, 0)	(2, 18)	135	11.5	54268–58470

Table C.3 Astrometric and rotational parameters for the 23 non-glitching pulsars in our sample. Asymmetric confidence intervals are individually listed. Lower- and upper-limits on $\dot{\nu}$, μ_α or μ_δ are given by two comma-separated values in parentheses. *Position and proper-motion fixed to values from very-long baseline interferometry (Deller et al., 2019).

PSRJ	RAJ (hh:mm:ss)	DECJ (°:′:″)	PEPOCH (MJD)	ν (Hz)	$\dot{\nu}$ (10^{-15} s $^{-2}$)	$\ddot{\nu}$ (10^{-24} s $^{-3}$)	μ_α (mas yr $^{-1}$)	μ_δ (mas yr $^{-1}$)	N_{ToA}	T (yr)	MJD range (MJD)
J0614+2229	06:14:17.0058(1)*	22:29:56.848(1)*	56000	2.9851428242(6)	-526.51(1)	(0, 5)	-0.23(5)*	-1.22(7)*	125	10.9	54505–58469
J0627+0706	06:27:44.17(4)	07:06:33.4(2)	55382	2.1013735312(3)	-13.155(3)	(-0.5, 0.1)	(-21, 8)	(14, 115)	126	10.7	54548–58470
J0659+1414	06:59:48.19(5)	14:14:21.3(3)	55586	2.59794880620(3)	-37.09365(8)	0.68(2)	49(13)	78^{+75}_{-68}	129	10.9	54505–58469
J0842-4851	08:42:05.33(5)	-48:51:20.65(3)	56022	1.55191806425(5)	-2.30338(6)	(-0.46, -0.14)	(-13, 15)	(-21, 9)	104	8.7	55363–58531
J0855-4644	08:55:36.16(2)	-46:44:13.46(1)	55288	15.4587515567(1)	-173.5507(4)	1.53(7)	(-7, 1)	(-3, 5)	118	11.6	54220–58469
J1003-4747	10:03:21.54(2)	-47:47:01.40(1)	55344	3.25654603313(2)	-2.19707(2)	0.04(2)	-12(2)	21(2)	199	12.3	53973–58469
J1057-5226	10:57:59.012 $^{+0.10}_{-0.09}$	-52:26:56.49(7)	55311	5.07321840599(8)	-15.01619(6)	(-0.28, -0.02)	49(4)	-6(5)	397	26.4	48814–58469
J1138-6207	11:38:21.7(3)	-62:07:59.0(1)	55355	8.5056937000(1)	-90.303(1)	(0.1, 0.4)	(-160, -40)	(20, 140)	89	6.8	54220–56708
J1327-6400	13:27:10(2)	-64:00:13.2(9)	55411	3.56265606571(9)	-40.327(2)	(-5, 11)	(-17, 442)	(61, 550)	86	6.8	54220–56708
J1359-6038	13:59:59.33(1)	-60:38:17.998 $^{+0.002}_{-0.001}$	55426	7.84268967507968(1)	-38.97175768(1)	(0.14, 0.46)	-4(4)	10(5)	808	28.9	47913–58469
J1541-5535	15:41:47(1)	-55:34:07.5 $^{+0.9}_{-1.0}$	55374	3.3799534494(6)	-84.0607(8)	(6, 13)	(-370, 10)	(-510, 70)	146	11.6	54220–58469
J1600-5044	16:00:53.03(3)	-50:44:21.0153 $^{+0.03}_{-0.04}$	57600	5.19197591041(8)	-13.64417(5)	(0.9, 0.39)	(-5, 3)	(-5, 11)	185	14.9	53040–58469
J1626-4807	16:26:42.5(7)	-48:07:56(1)	55292	3.4021201284(3)	-20.2284(9)	(-9, 11)	-	-	94	8.0	54220–57129
J1627-4706	16:27:28.79 $^{+0.08}_{-0.07}$	-47:06:49.3(1)	53165	7.1050074462(2)	-8.7355(2)	(-0.14, -0.02)	(-7, 33)	(-40, 60)	134	10.7	52807–56708
J1701-4533	17:01:29.13(2)	-45:33:49.18(4)	48360	3.0968518767(3)	-0.49801(5)	(-5.5, -0.7)	-	-	56	5.1	56682–58531
J1716-4005	17:16:42.0(5)	-40:05:27(1)	54942	3.207052805(4)	-2.921(1)	(-13, -5)	-	-	38	3.4	57276–58531
J1721-3532	17:21:32.76(7)	-35:32:48.3(3)	55456	3.56591744853(3)	-32.0185(1)	(0.22, 0.34)	(-14, 10)	(-60, 60)	252	18.0	51879–58470
J1726-3530	17:26:07.6(1)	-35:29:58(4)	55432	0.900460470(1)	-98.25(3)	20(2)	(-300, 60)	-	201	16.5	50681–56709
J1750-3157	17:50:47.31(2)	-31:57:44.3(1)	50271	1.09846293993(2)	-0.023700(4)	(-1, -0.8)	-	-	42	12.4	53974–58500
J1835-0643	18:35:05.56(8)	-06:43:06.9(3)	55365	3.2695714904(2)	-43.2078(2)	(1.2, 2.0)	(-22, 24)	(-120, 40)	113	11.3	54268–58404
J1841-0425	18:41:05.7(1)	-04:25:20.3(5)	55402	5.3720208212(3)	-18.459(1)	1.2(1)	(-21, 6)	(0, 50)	139	11.5	54268–58470
J1844-0256	18:44:30.1(3)	-02:54:09.7 $^{+0.8}_{-0.7}$	54919	3.66337798586(4)	-38.3135(2)	(-0.7, 0.1)	-	-	61	4.3	54268–55822